Observers Are All You Need

B. Müller Peter Nguyen

April 1, 2026

Paper release: `r1367` **Released:** April 1, 2026

Abstract

Observer-Patch Holography (OPH) asks whether familiar physics can be rebuilt from a very small starting point: no observer sees the whole world at once, each observer has access only to local data, and neighboring descriptions must agree where they overlap. In that picture, spacetime, gauge structure, particles, and classical records are not put in by hand at the start. They are features of whatever survives global consistency among many partial viewpoints on a finite-capacity screen.

This paper is the synthesis surface for the OPH suite. It pulls together the fixed-cutoff collar and higher-gauge theorem packages, the consensus/fixed-point formulation of overlap repair, the conditional route to Lorentzian geometry and Einstein dynamics, the compact-gauge route to the realized Standard Model quotient $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)/\mathbb{Z}_6$, the particle-spectrum lane, and the screen microphysics and observer lane. It shows, in one place, what is closed, what is conditional, and what sits on an explicit missing object.

At fixed cutoff, OPH has a sharp local picture: patches glue, higher-gauge defects are controlled, physical UV uniqueness is quotient-level and ancilla-stable, and the microphysics lane has an explicit regulated architecture together with closed measurement and checkpoint/restoration packages. On the continuum side, OPH gives a conditional Lorentzian/Jacobson-type gravity branch and a conditional compact-gauge route whose realized low-energy output is the Standard Model quotient above with the exact hypercharge lattice and the counting chain $N_g = 3$ then $N_c = 3$. On the particle side, the structural massless carriers are fixed, the electroweak lane is closed at theorem level, the Higgs/top stage is emitted on its secondary branch, and the neutrino lane matches the observed mixing/hierarchy pattern while leaving one reduced normalization invariant open; the charged-lepton, quark, and hadron lanes keep their remaining closure objects explicit. The biggest open front is the continuum/BW lift: an explicit microscopic refinement whose scaling limit lands in the geometric modular phase.

Contents

1 Introduction

This synthesis paper reads the OPH suite on one synchronized theorem and status surface. It brings together the recovered core from Ref. [1], the finite patch-net consensus spine from Ref. [2], the screen-microphysics engineering lane from Ref. [3], and the downstream particle-spectrum lane from Ref. [4]. The guiding question is whether overlap consistency for local observer descriptions on a finite-capacity screen is strong enough to recover the structural features of low-energy physics, and if so, how much of that recovery is theorem-level and how much depends on explicit scaling-limit, categorical, or continuation premises.

The answer is sharply tiered. Fixed-cutoff collar structure, higher-gauge defect transport, the consensus/fixed-point formulation, the quotient-level UV uniqueness statement, the Born-rule record package, and the checkpoint/restoration package are theorem-bearing on their stated finite surfaces. The Lorentz/null-modular/Einstein chain and compact gauge reconstruction remain conditional continuum or categorical branches. The particle lane is the most status-split part of the suite: the structural carrier skeleton and realized Standard Model branch are fixed, the electroweak stage is closed, the Higgs/top stage is a secondary quantitative branch, the neutrino lane is shape-accurate with one reduced normalization invariant open, and the charged-lepton, quark, and hadron lanes remain explicitly unfinished for different reasons.

The paper is organized by lane. Section 2 gives the synthesis-level map. Section 3 carries the recovered-core trunk: axioms, premise ledger, collar tools, gluing, the Lorentz/gravity branch, and the compact-gauge/Standard-Model branch. Section 4 records the consensus lane. Section 5 summarizes the particle lane. Section 6 records the screen-microphysics and observer lane. Section 7 keeps the worldsheet lane at its conditional status. Section 8 gives the cross-lane status and open fronts. The appendices collect extra architecture and interpretive material that is useful in the synthesis paper but not part of the recovered core itself.

2 Overview of Results

The OPH suite currently supports the following lane-level split.

1. **At fixed cutoff, local structure and repair are sharp.** Technically, the collar theorem package is explicit and finite-range, the genuinely noncentral higher-gauge branch is closed at fixed cutoff, and the consensus lane gives a unique schedule-independent normal form on the gauge quotient under the stated finite patch-net hypotheses.

2. **Relativity appears as a controlled continuum branch.** Technically, the Lorentz/null-modular/Einstein chain is a refinement/scaling-limit branch with explicit canonical technical premises.

3. **The Standard Model gauge skeleton appears as a constrained consistency output.** Technically, compact gauge reconstruction and the realized $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)/\mathbb{Z}_6$ branch follow under the stated transportability and categorical premises, with the realized counting chain $N_g = 3$ then $N_c = 3$.

4. **The particle lane has real outputs and a nonuniform closure profile.** Technically, the structural carriers are exact, the electroweak stage is closed, the Higgs/top stage is quantitatively emitted on a secondary branch, and the charged-lepton, quark, neutrino, and hadron lanes carry named remaining closure objects.

5. **The screen-microphysics lane is explicit and concrete.** Technically, it contains an explicit regulated reference architecture, sharpened patch-net and edge-sector branches, a closed fixed-cutoff Born-rule package, and a closed checkpoint/restoration package.
6. **The UV picture is sharper at fixed cutoff than in the continuum.** Technically, OPH fixes the physical UV branch only modulo implementation hiding and inert ancillary refinement; the UV burden is the continuum BW/geometric lift and the refinement-limit transportable-sector lift.
7. **The string/worldsheet lane remains conditional.** Technically, it is carried only as a continuation of the OPH heat-kernel edge-sector theorem, not as part of the recovered core.

Companion Papers

The detailed depth surfaces used throughout this synthesis are:

1. *A Conditional Reconstruction Program for Relativity and Standard Model Structure from Observer-Overlap Consistency: Observers Are All You Need (Compact)* [1], which carries the premise ledger, the recovered-core Lorentz/gravity route, and the compact-gauge/Standard-Model branch.
2. *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics* [2], which carries the finite patch-net fixed-point, defect, quotient, and record-layer consensus package.
3. *Screen Microphysics, Patches, and Observer Synchronization in OPH* [3], which carries the regulated reference architecture together with the measurement and observer checkpoint/restoration packages.
4. *Deriving the Particle Zoo from Observer Consistency* [4], which carries the structural-to-family particle route and the per-sector status ledger.

3 Recovered Core: Foundations and Structural Branches

OPH is an observer-centric reconstruction program for fundamental physics. Physical data live on a horizon screen S^2, and reality is encoded by the mutual consistency of overlapping patch descriptions. This paper works with five OPH axioms, two explicit external inputs used only in the quantitative branches, and theorem-local technical premises stated where needed for the scaling, gauge, and transport arguments below.

3.1 Overview

The recovered core of the paper is the D1–D5 and D7–D9 chain: overlap-consistency normal form, collar and edge-center structure, the conditional Lorentz and Einstein branches, compact gauge reconstruction, and the realized Standard Model branch with exact hypercharge structure, $N_g = 3$, and $N_c = 3$. Here and throughout the synthesis, labels such as D1–D12 are internal documentary node labels for the OPH derivation ledger; the local key appears in Section 3.2.6, the SM/GR derivation paper *Recovering Relativity and Standard Model Structure from Observer-Overlap Consistency* gives the same ledger in shorter form [1], and the public paper ledger has no separate

D11 row. D6 is a separate input-dependent corollary tied to total screen capacity. D12 collects phenomenological continuations kept distinct from the recovered core.

The Lorentz/null-modular/Einstein chain is a scaling-limit statement with explicit fixed-cutoff control. The gauge chain reconstructs a compact group from a refinement-stable sector colimit and then specializes to the realized MAR-admissible branch. The worldsheet/string material and the flavor, neutrino, hadron, dark-sector, and spectroscopy branches are kept separate from the recovered core.

Two external continuous inputs appear only where declared. The pixel area $P \equiv a_{\mathrm{cell}}/\ell_P^2$ enters deferred quantitative branches, and the screen capacity $N_{\mathrm{scr}} \equiv \log \dim \mathcal{H}_{\mathrm{tot}}$ enters the cosmological-capacity branch. They are implementation inputs, not additional axioms.

3.2 Model and Axioms

3.2.1 Observers and access model

An observer O is a tuple $(P_O, \mathcal{A}(P_O), \rho_O, R_O)$ where:

- $P_O \subset S^2$ is a connected screen patch (the observer's access region).
- $\mathcal{A}(P_O)$ is the von Neumann algebra associated to P_O.
- ρ_O is the local state, obtained by restricting the global state to $\mathcal{A}(P_O)$.
- R_O is a set of records: stable internal correlations within P_O.

Observers are internal patterns in the global state. Different observers correspond to different patches and their compatible marginals.

3.2.2 Screen, patches, and algebra net

We work in a single static patch with a horizon screen S^2. Each connected subregion $P \subset S^2$ is assigned a von Neumann algebra $\mathcal{A}(P)$. The net satisfies isotony:

$$P \subset Q \implies \mathcal{A}(P) \subset \mathcal{A}(Q).$$

A global state ω is a positive linear functional on the inductive-limit algebra. Overlap consistency is imposed algebraically: for overlaps $P_1 \cap P_2$, ω restricted to $\mathcal{A}(P_1 \cap P_2)$ is the same from either side.

3.2.3 Five OPH axioms

This paper uses five OPH axioms. The only labeled technical premises used below are the canonical premises **T1**–**T7** stated in Section 1.4. Other theorem-local hypotheses are written out in prose where needed rather than assigned a separate label family.

1. Screen Net. A horizon screen S^2 carries a net of algebras $P \mapsto \mathcal{A}(P)$.

2. Overlap Consistency. Local states agree on shared observables for any overlap.

3. Local MaxEnt and Refinement Stability. At the UV scale the realized state maximizes entropy subject to a finite family of gauge-invariant local constraints, and the resulting family of states is stable under coarse-graining so symmetry-allowed relevant operators are not held at zero by unexplained fine tuning.

4. Recoverable Generalized Entropy. A generalized entropy functional exists on caps, obeys quantum focusing on null generators, and comes with the recoverability structure used throughout the manuscript: collar tripartitions have small conditional mutual information with controlled recovery maps, with stronger collar/null-strip hypotheses stated explicitly where needed.

5. Minimal Admissible Realization (MAR). On the admissible low-energy branch, the realized sector package is the lexicographically minimal one under the complexity vector $C(\mathfrak{S})$. MAR is an explicit structural-economy axiom on admissible realized branches, not a theorem derived from the preceding four axioms.

3.2.4 External inputs and theorem-local technical premises

The quantitative branches discussed here use exactly two external continuous inputs:

$$P \equiv a_{\text{cell}}/\ell_P^2 = 1.63094, \qquad N_{\text{scr}} \equiv \log\dim\mathcal{H}_{\text{tot}} \sim 10^{122}.$$

Here P is calibrated from the gauge sector and N_{scr} is inferred from the observed cosmological constant. The first feeds deferred quantitative branches; the second enters the cosmological-capacity branch and a separate order-of-magnitude neutrino estimate. These are implementation inputs. They are not extra axioms.

Later theorem statements use only the canonical technical-premise labels **T1**–**T7**. Fixed-cutoff regulator bookkeeping, quasi-local propagation, endpoint control, collar decomposition, and related branch-internal statements are cited directly from axioms and earlier theorems rather than through a second premise-label system.

Local MaxEnt at the regulator scale. At the regulator scale ℓ_{UV}, the global state ω maximizes von Neumann entropy subject to:

1. A finite set $\{O_a\}$ of gauge-invariant local operators, each supported on a ball of radius $\leq r_0 = O(\ell_{\text{UV}})$.
2. Constraint equations $\langle O_a(x)\rangle = c_a$ for each cell x in the UV lattice.
3. Optionally, a finite number of global constraints (total energy, charge).

This is the minimal specification needed to derive the local Gibbs form of Lemma 2.6 from Axiom 3.

Clarification (MaxEnt $\neq$ thermal equilibrium). MaxEnt here is **local state selection** given constraints, not "the universe is in thermal equilibrium." The Lagrange multipliers (inverse temperatures) may vary slowly in space and time. Non-equilibrium physics appears as gradients in these multipliers and as controlled violations of exact Markov additivity under the explicit mixing hypotheses stated later in Section 2.3. Equilibrium is an approximation regime with explicit error terms.

Rotationally invariant constraint branch. Constraint sets are SO(3)-invariant on S^2.

Gauge as overlap redundancy at fixed cutoff. Overlap identifications are not unique; the freedom that leaves overlap observables invariant forms a local groupoid.

Central-defect subbranch. On triple overlaps, when the failure of strict coherence is central, one has

$$\varphi_{ij}\varphi_{jk}\varphi_{ki} = \text{Ad}(z_{ijk}), \qquad z_{ijk} \in Z(\mathcal{A}_{ijk}).$$

Collar double-scaling hypothesis. There exists a UV length ℓ_{UV} such that for any cap C and collar width δ, in the refinement limit $\delta \to 0$ and $\ell_{\text{UV}} \to 0$ with $\delta/\ell_{\text{UV}} \to \infty$, the Markov error satisfies

$$I(A_\delta : D_\delta \mid B_\delta)_\omega \leq \varepsilon(\delta/\ell_{\text{UV}}), \qquad \varepsilon(x) \to 0 \text{ as } x \to \infty.$$

See Section 2.3 for the collar tripartition definitions and the regulated EC proof that yields this limit from the fixed-cutoff realized presentation.

Bisognano–Wichmann geometric branch. The refinement-stable OPH branch for caps lies in the geometric modular class, together with the controlled collar/refinement limit that carries the Markov and recovery remainders explicitly. In the canonical numbered list below, this is the content of **T1** together with the scaling-limit scope clause **T3**.

Derived quasi-local propagation and modular-locality control. At scale ℓ_{UV}, the dynamics generated by the effective Hamiltonian H_{eff} (from Lemma 2.6) has a finite Lieb-Robinson velocity v_{LR}: for local operators A, B supported on regions separated by distance d,

$$\|[A(t), B]\| \leq c\|A\|\|B\| \min(|R_A|, |R_B|) e^{-(d-v_{\mathrm{LR}}|t|)/\xi}$$

for $d > v_{\mathrm{LR}}|t|$, where $\xi = O(\ell_{\mathrm{UV}})$.

This is the branch-internal control statement that turns the quasi-local structure of the local-Gibbs branch into explicit support control for time-evolved operators. It is not an extra premise beyond the third OPH axiom.

> **Approximate modular covariance on the geometric branch.** Under the fixed-cutoff realized presentation, Lemma 2.6 (local Gibbs), the derived quasi-local propagation bound above, and the collar double-scaling plus mixing hypotheses of Section 2.3, the modular flow $\sigma_t^{\omega,C}$ maps $\mathcal{A}(R)$ into a slightly thickened region algebra:
>
> $$\sigma_t^{\omega,C}(\mathcal{A}(R)) \subseteq \mathcal{A}(R^{+v_{\mathrm{mod}}|t|}) \quad \text{up to error } \eta(d - v_{\mathrm{mod}}|t|),$$
>
> where R^{+s} denotes the s-neighborhood thickening, v_{mod} is a "modular propagation velocity" controlled by v_{LR} and local norm bounds, and d is the distance from R to ∂C.
>
> **Heuristic motivation.** The modular Hamiltonian $K_C = -\log \rho_C$ is quasi-local by Lemma 4.1a-b (modular additivity localizes it to the collar). The derived Lieb-Robinson control then bounds support spreading under $e^{iK_C t}$. In the double-scaling collar limit ($\delta/\ell_{\mathrm{UV}} \to \infty$), the thickening vanishes in macroscopic units. This is heuristic support for the tangent-limit package used later.
>
> **Continuum-limit heuristic.** Define the induced region flow
>
> $$f_t^C(R) := \lim_{\ell_{\mathrm{UV}} \to 0} R^{+v_{\mathrm{mod}}|t|}.$$
>
> Then $\sigma_t^{\omega,C}(\mathcal{A}(R)) = \mathcal{A}(f_t^C(R))$ becomes exact in the continuum limit, with error controlled by
>
> $$\eta(\delta) \lesssim 2\sqrt{\ln 2 \cdot c \cdot |\partial C|_{\mathrm{UV}}}\, e^{-\delta/(2\xi)}.$$

The discussion above is heuristic support for the tangent-limit package used in Theorem 4.2.

Refinement-stable MaxEnt branch. There exists a family of coarse-graining channels $\Phi_{\ell\to L}$ between UV scale ℓ and IR scale L such that the MaxEnt-selected states are self-similar under refinement,

$$\Phi_{\ell\to L}(\omega_\ell) = \omega_L,$$

with the constraint set fixed and finite. Equivalently, the MaxEnt family is an RG fixed point or a low-dimensional stable manifold determined only by the constraints.

Fixed-cutoff realized presentation. At a UV scale ℓ_{UV}, local patch algebras are type-I with finite-dimensional Hilbert spaces, and gauge-as-gluing is realized as a boundary group action whose fixed-point algebra defines physical observables. Section 2.3 proves the resulting fixed-cutoff collar package directly from overlap consistency on this realized presentation.

External mathematical inputs: SSA and recovery theorems (Petz 1986, 1988; Fawzi and Renner 2015), Jacobson's entanglement-equilibrium derivation (Jacobson 1995, 2016), and one of the following EFT bridges: (i) the null-surface modular route (Section 5.2), or (ii) a UV CFT regime on sufficiently small caps (Section 5.3). For SM contact we also use the Doplicher-Roberts reconstruction (Doplicher and Roberts 1989, 1990) once localized transportable sectors are assumed in the small-region limit. Full citations appear in the References.

Canonical technical premises. The only labeled technical premises used in this manuscript are:

- **T1**: the branch condition that the refinement-stable cap net lies in the Bisognano–Wichmann geometric modular class for caps and their conformal images.
- **T2**: vanishing of the relevant gluing obstruction class whenever global sector transportability is invoked, written $[z] = 0$ on the central branch or $q_\Sigma = 0$ on the genuinely noncentral branch.
- **T3**: all Lorentz/null-modular/Einstein statements are claims about a controlled refinement/scaling limit of the regulator net, not literal exactness at fixed finite cutoff.
- **T4**: fixed-cap generalized-entropy stationarity for the admissible first-variation class used in the Jacobson branch.
- **T5**: symmetric braiding in the $3+1$D EFT branch.
- **T6**: a bosonic Tannakian fiber-functor premise, or else an explicit super-Tannakian fork.
- **T7**: whenever compact gauge reconstruction is invoked, the sector category means the directed colimit of transportable edge sectors with objectwise finite-dimensional fibers.

3.2.5 Notation

- ρ_C: reduced state on cap C.
- $K_C := -\log \rho_C^\omega$: modular Hamiltonian of the reference state.
- B_C: geometric generator of the cap-preserving conformal dilation.
- $S_{\mathrm{gen}}(C)$: generalized entropy on a cap.
- ℓ_{UV}: UV length scale of the refined screen net.
- δ: collar width around a cap boundary.

3.2.6 Summary: premise ledger and dependency DAG

For reference, the paper uses the five OPH axioms above, the two external inputs P and N_{scr} only in the branches that require them, and the theorem-local technical premises declared in Section 1.4.

Item	Meaning, typical use, and status
Axiom 1	Screen Net; use: global net structure; status: axiom.
Axiom 2	Overlap Consistency; use: compatibility on overlaps; status: axiom.
Axiom 3	Local MaxEnt and Refinement Stability; use: state selection and RG stability; status: axiom.
Axiom 4	Recoverable Generalized Entropy; use: focusing plus recoverability structure; status: axiom.
Axiom 5	Minimal Admissible Realization (MAR); use: selection on admissible gauge branches; status: axiom.
P	Pixel area $a_{\mathrm{cell}}/\ell_P^2$; use: deferred quantitative calibration input; status: external input.
N_{scr}	Total screen capacity; use: cosmological-capacity branch; status: external input.

Item	Meaning, typical use, and status
T1	Bisognano–Wichmann geometric branch for caps and their conformal images; use: Theorem 4.2 and the Lorentz branch; status: canonical technical premise.
T2	Vanishing relevant transport obstruction; use: gauge reconstruction and realized product-group branch; status: canonical technical premise when invoked.
T3	Controlled refinement/scaling-limit scope clause; use: Lorentz/null-modular/Einstein branch; status: canonical technical premise.
T4	Fixed-cap generalized-entropy stationarity; use: Jacobson branch; status: canonical technical premise.
T5	Symmetric braiding in the $3+1$D EFT branch; use: compact gauge reconstruction; status: canonical technical premise.
T6	Bosonic Tannakian fiber functor or explicit super-Tannakian fork; use: compact gauge reconstruction; status: canonical technical premise.
T7	Directed colimit of transportable edge sectors with objectwise finite-dimensional fibers; use: compact gauge reconstruction; status: canonical technical premise.

Master dependency map. The completion program is organized into the following nodes. The point of the table is to make explicit which results use the OPH axioms alone, which import standard mathematics, and which require extra physical premises or external inputs.

Node	Main-manuscript output	Immediate OPH ingredients	Standard mathematics used	Extra physical premises / external inputs	Claim tier
D1	Theorem 3.1: schedule-independent normal-form statement	overlap-consistency setup on finite patch nets	Newman's lemma	finite patch net, termination, and local confluence	structural theorem
D2	Section 2.3 and §5.4: EC/Markov collar and generalized-entropy split	Axioms 1–4 plus the derived fixed-cutoff regulator/collar package when invoked	HJPW; Petz; Fawzi–Renner	exact-Markov or idealized recoverability hypotheses when exact identities are used; coarse-grained $A/(4G)$ dictionary	structural lemma/proposition layer
D3	Theorems 4.2–4.3: geometric modular flow and Lorentz branch	Axioms 1–4 together with SO(3)-invariant constraint data	Bisognano–Wichmann modular geometry; conformal classification on S^2	T1+T3	conditional scaling-limit theorem
D4	Section 5.2: null modular bridge	D2+D3	Borchers–Wiesbrock and endpoint differentiation	the theorem-local inherited-strip decomposition and exact-or-controlled Markov hypotheses of Section 5.2, plus the downstream density-upgrade hypotheses and the relativistic null-stress premise	conditional bridge theorem
D5	Theorem 5.1: Einstein branch	D3+D4	Jacobson small-ball mathematics; null-to-tensor reconstruction modulo a metric term	T4 together with the locally Lorentzian $d = 4$ scaling regime	conditional scaling-limit theorem
D6	capacity/Λ reduction statements	D5 leaves $+\Lambda g_{ab}$ undetermined	de Sitter entropy relation and dimensional analysis	external input N_{scr} and the screen-capacity identification	input-dependent corollary / estimate
D7	Theorem 6.1: compact gauge reconstruction	Axioms 1–4	Doplicher–Roberts / Tannaka reconstruction	T2 and T5–T7	conditional structural theorem

Node	Main-manuscript output	Immediate OPH ingredients	Standard mathematics used	Extra physical premises / external inputs	Claim tier
D8	Section 6.2: product gauge structure up to finite quotient	D7 + Axiom 5 (MAR)	compact Lie representation classification; Schur's lemma	connected admissible class, one connected abelian factor, and vanishing relevant transport obstruction	realized-branch theorem
D9	Theorem 6.13, Proposition 6.9, Theorem 6.14, Proposition 6.6: hypercharge lattice, the counting chain $N_g = 3$ then $N_c = 3$, and the $\mathbb{Z}_6$ quotient	D8	anomaly algebra; Witten anomaly; CKM CP counting	realized one-generation/one-Higgs package and MAR admissibility premises	realized-branch theorem/corollary chain
D12	downstream flavor, dark-sector, baryogenesis, spectroscopy, and string continuations	various subsets of D6 and D9	branch-specific EFT and phenomenological manipulations	explicit additional ansätze	phenomenological continuation / program branch

In particular, D3–D5 are scaling-limit branches rather than fixed-cutoff matrix-algebra theorems; D7–D9 are realized-branch structural outputs in the bosonic internal-gauge sector; D6 is input-dependent; and D12 remains phenomenological.

3.3 Information-Theoretic Tools

3.3.1 Strong subadditivity and Markov states

For any tripartite state ρ_{ABC},

$$I(A:C \mid B) := S(AB) + S(BC) - S(B) - S(ABC) \geq 0.$$

Exact Markov states satisfy $I(A:C \mid B) = 0$ and admit a recovery map:

$$\rho_{ABC} = (\mathrm{id}_A \otimes \mathcal{R}_{B\to BC})(\rho_{AB}).$$

3.3.2 Approximate recovery

If $I(A:C \mid B) \leq \varepsilon$ (bits), there exists a CPTP recovery map $\mathcal{R}$ with

$$\|\rho_{ABC} - (\mathrm{id}_A \otimes \mathcal{R})(\rho_{AB})\|_1 \leq 2\sqrt{\ln 2\,\varepsilon}.$$

3.3.3 Collar refinement and sufficient mechanisms

Fix a cap $C \subset S^2$ with boundary circle. For collar width δ define

$$B_\delta := \{x \in S^2 : \mathrm{dist}(x, \partial C) \leq \delta\}, \qquad A_\delta := C \setminus B_\delta, \qquad D_\delta := (S^2 \setminus C) \setminus B_\delta.$$

Then $S^2 = A_\delta \cup B_\delta \cup D_\delta$ with A_δ and D_δ interacting only through B_δ. The collar double-scaling hypothesis is the requirement that $I(A_\delta : D_\delta \mid B_\delta) \to 0$ in the refinement limit. At fixed regulator scale, the same finite-dimensional realized presentation supplies the patch-net and overlap-gluing data used below. We therefore isolate that realized presentation first, then separate the exact-Markov and quantitative routes.

Derived regulator realization. Choose a finite UV cellulation at scale ℓ_{UV} and let R be a finite union of cells. Hilbertize each cell's finite local data to a finite-dimensional space $\tilde{\mathcal{H}}_i \cong \mathbb{C}^{n_i}$. The extended algebra before quotienting by overlap redundancy is

$$\widetilde{\mathcal{A}}(R) = \mathcal{B}(\tilde{\mathcal{H}}_R), \qquad \tilde{\mathcal{H}}_R = \bigotimes_{i\subset R} \tilde{\mathcal{H}}_i.$$

Overlap-preserving changes of local trivialization act only on cut data. On each finite-dimensional chart, their unitary image has compact closure, so one may represent the boundary gluing redundancy by a compact group $G_{\partial R} \subset U(\tilde{\mathcal{H}}_R)$. The physical algebra is the fixed-point algebra

$$A(R) = \widetilde{\mathcal{A}}(R)^{G_{\partial R}} = \mathcal{B}(\tilde{\mathcal{H}}_R)^{G_{\partial R}}.$$

This is the fixed-cutoff realized presentation used throughout the collar analysis.

Theorem 2.3 (EC from regulated overlap gluing). For a collar B_δ around a cap boundary Σ, there is a canonical decomposition

$$H_{B_\delta} = (\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{G_\Sigma} = \bigoplus_\alpha (H_{b_L^\alpha} \otimes H_{b_R^\alpha}),$$

with

$$Z(A(B_\delta)) = \bigoplus_\alpha \mathbb{C}\,\mathbf{1}_\alpha,$$

such that $\mathcal{A}(A_\delta B_\delta)$ acts only on $H_{b_L^\alpha}$ and $\mathcal{A}(B_\delta D_\delta)$ acts only on $H_{b_R^\alpha}$ within each block.

Proof. Split the collar into half-collars B_L and B_R meeting on $\Sigma = \partial C$. By the realized regulator presentation above, the physical collar Hilbert space is the diagonal invariant subspace $(\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{G_\Sigma}$. Decompose each side into irreps:

$$\tilde{\mathcal{H}}_{B_L} = \bigoplus_\alpha (V_\alpha \otimes H_{b_L^\alpha}), \qquad \tilde{\mathcal{H}}_{B_R} = \bigoplus_\beta (V_\beta^* \otimes H_{b_R^\beta}).$$

Then

$$\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R} = \bigoplus_{\alpha,\beta} (V_\alpha \otimes V_\beta^*) \otimes (H_{b_L^\alpha} \otimes H_{b_R^\beta}).$$

By Schur's lemma,

$$(V_\alpha \otimes V_\beta^*)^{G_\Sigma} \cong \begin{cases} \mathbb{C}, & \alpha = \beta, \\ 0, & \alpha \neq \beta. \end{cases}$$

Therefore the invariant subspace is

$$H_{B_\delta} = \bigoplus_\alpha (H_{b_L^\alpha} \otimes H_{b_R^\alpha}),$$

as claimed. The invariant algebra is

$$A(B_\delta) = \bigoplus_\alpha (B(H_{b_L^\alpha}) \otimes B(H_{b_R^\alpha})),$$

so the center is generated by the block projectors. Adjacent region algebras act on the left or right factor only because the gluing action is supported on Σ. QED.

Remark. On the central-defect subbranch, replace G_Σ by its central extension. The sector label α then ranges over irreps of the extension; the decomposition is unchanged.

We refer to the decomposition in Theorem 2.3 as **edge-center completion (EC)**.

Exact Markov route. If the reference state on $A_\delta B_\delta D_\delta$ is exact Markov, or in an explicitly stated idealized limit that reduces to exact Markovity, then

$$\rho_{A_\delta B_\delta D_\delta} = \bigoplus_\alpha p_\alpha \big(\rho_{A_\delta b_L^\alpha} \otimes \rho_{b_R^\alpha D_\delta}\big), \qquad I_\omega(A_\delta : D_\delta \mid B_\delta) = 0.$$

This is the HJPW normal form applied to the EC blocks, and it is the exact identity used for literal Markov-modular equalities.

Interpreting collar refinement as the inductive limit of these regulators with $\delta/\ell_{\mathrm{UV}} \to \infty$, Theorem 2.3 supplies the kinematic edge-center decomposition. Exact Markovity is one idealized route. The following lemma and axiom provide the quantitative decay route used when the manuscript keeps the approximation explicit.

Lemma 2.6 (MaxEnt with local constraints implies local Gibbs form). Under the local MaxEnt branch of Axiom 3 and the finite-dimensional regulator realization above, the MaxEnt state has the Gibbs form

$$\omega = \frac{e^{-H_{\mathrm{eff}}}}{\operatorname{Tr} e^{-H_{\mathrm{eff}}}}, \qquad H_{\mathrm{eff}} = \sum_x \sum_a \lambda_a O_a(x) + (\text{global terms}),$$

where the sum runs over UV cells x and constraint operators O_a. The effective Hamiltonian H_{eff} is quasi-local with range $O(\ell_{\mathrm{UV}})$.

Proof. On a finite-dimensional algebra, the unique state maximizing $S(\rho) = -\operatorname{Tr}(\rho \log \rho)$ subject to linear constraints $\operatorname{Tr}(\rho O_i) = c_i$ is given by Lagrange multipliers:

$$\rho = \frac{e^{-\sum_i \lambda_i O_i}}{\operatorname{Tr} e^{-\sum_i \lambda_i O_i}}.$$

Strict concavity of von Neumann entropy ensures uniqueness. When the constraints are "translated local" (the same O_a at each cell x), the exponent is a sum of local terms. QED.

Exponential mixing hypothesis. There exist constants c and correlation length $\xi = O(\ell_{\mathrm{UV}})$ such that

$$I_\omega(A_\delta : D_\delta \mid B_\delta) \le c\,|\partial C|_{\mathrm{UV}}\, e^{-\delta/\xi}, \qquad |\partial C|_{\mathrm{UV}} \sim \frac{\operatorname{length}(\partial C)}{\ell_{\mathrm{UV}}}.$$

This is the standard clustering/mixing condition for local Gibbs states, equivalent to assuming that the MaxEnt state lies in a Dobrushin uniqueness regime or has a uniform spectral gap. It is a natural UV-state condition and is not derived from Axiom 3 alone.

Theorem 2.5 (Local Gibbs + mixing implies collar refinement). Under Lemma 2.6 and the exponential mixing hypothesis above, the collar double-scaling hypothesis holds in the limit $\delta \to 0$, $\ell_{\mathrm{UV}} \to 0$ with $\delta/\ell_{\mathrm{UV}} \to \infty$.

Proof. The mixing bound above has polynomial growth in $|\partial C|_{\mathrm{UV}}$ and exponential decay in $\delta/\ell_{\mathrm{UV}}$. In the double-scaling limit the exponential dominates, so $I_\omega(A_\delta : D_\delta \mid B_\delta) \to 0$. QED.

This bound is the quantitative hinge for constructive gluing.

3.3.4 Concrete UV realization: quantum link models

Section 2.3 internalizes the regulator package: the fixed-cutoff type-I algebra and boundary fixed-point structure are the realized form of the screen net plus overlap gluing. Quantum link models are included here only as an explicit microscopic example of that structure, not as a separate axiom layer.

UV regulator. Triangulate S^2 at scale ℓ_{UV}, giving vertices v, oriented links ℓ, and plaquettes p. Refinement corresponds to $\ell_{\mathrm{UV}} \to 0$ with increasing lattice size.

Degrees of freedom. Attach to every oriented link ℓ a **finite-dimensional** Hilbert space $\mathcal{H}_\ell$. In ordinary Wilson lattice gauge theory, the continuum/refinement-limit edge description is modeled by $L^2(G)$ (infinite-dimensional for continuous G), but the microscopic OPH regulator is instead a **quantum link model** with finite-dimensional link Hilbert spaces that preserve gauge symmetry in operator form [32]. Optionally attach matter Hilbert spaces $\mathcal{H}_v$ at vertices. Then:

$$\tilde{\mathcal{H}}_{\mathrm{total}} = \bigotimes_\ell \mathcal{H}_\ell \otimes \bigotimes_v \mathcal{H}_v,$$

finite-dimensional on any finite lattice. This is a concrete realization of the extended type-I presentation used in Section 2.3.

Boundary gluing as Gauss-law invariants. Define a local gauge transformation group G_v at each vertex v acting on incident links (and matter at v). Physical states satisfy:

$$|\psi\rangle \in \mathcal{H}_{\mathrm{phys}} \quad \Longleftrightarrow \quad U(g_v)|\psi\rangle = |\psi\rangle \;\; \forall\, v,\, g_v \in G_v.$$

Equivalently: $\mathcal{H}_{\mathrm{phys}} = \tilde{\mathcal{H}}_{\mathrm{total}}^{\prod_v G_v}$.

Region algebras. For any region $R \subset S^2$, define an extended Hilbert space $\tilde{\mathcal{H}}_R$ from the links/vertices in R. The **boundary gauge group** $G_{\partial R}$ acts on the cut degrees of freedom (the "half-links" ending on ∂R). Define:

$$\mathcal{A}(R) = \mathcal{B}(\tilde{\mathcal{H}}_R)^{G_{\partial R}}.$$

This is the concrete quantum-link instance of the fixed-point algebra used in Section 2.3. The same definition gives isotony, overlap consistency, and the edge-center structure on collars.

Why EC follows immediately. Take a cap C and a collar B_δ around ∂C. Because the *only* coupling between inside and outside is through the boundary gauge constraint, the collar Hilbert space decomposes into superselection blocks labeled by boundary irreps:

$$\mathcal{H}_{B_\delta} \cong \bigoplus_\alpha (H_{b_L^\alpha} \otimes H_{b_R^\alpha}),$$

with center generated by the projectors P_α. This is precisely the Schur-lemma mechanism of Theorem 2.3. The labels α are the familiar edge-mode / electric-flux labels appearing whenever one factorizes gauge theories across an entangling cut [33]. Exact Markovity depends on the state and is not forced by the decomposition alone.

Dynamics and MaxEnt. The natural Hamiltonian is a 2+1D lattice gauge Hamiltonian on the screen worldvolume: plaquette ("magnetic") terms, electric terms on links, vertex Gauss terms as constraints, plus local matter couplings. In quantum link form this remains finite-dimensional per link while behaving like gauge theory in the continuum limit. Then the MaxEnt assumption becomes concrete: the MaxEnt state is a Gibbs state $\rho \propto e^{-\sum_i \lambda_i O_i}$ with quasi-local O_i, precisely the local-Gibbs regime.

Geometry and G. This microphysics naturally supplies the emergent geometric objects:

- **Edge entropy / area operator:** $L_C = \sum_\alpha (\log d_\alpha) P_\alpha$ becomes "log of boundary irrep dimension" in the gauge link model.
- **Newton constant G:** the conversion factor between edge entropy density per boundary UV cell and macroscopic geometric area.

Thus area is an operator living in the center of the boundary algebra, because in gauge systems the center is where the cut labels live.

What this example does and does not close. The quantum-link realization makes the fixed-cutoff matrix/fixed-point bookkeeping concrete. What it does **not** automatically guarantee is that modular flow on caps becomes geometric conformal dilation with the 2π KMS normalization, that is, the continuum modular-covariance regime encoded by **T1** and **T3**. That requires the state to sit in a regime that is effectively relativistic/QFT-like in the continuum limit. Viable architectures for this include holographic quantum error-correcting codes [34] and quantum double / string-net Hamiltonians [35].

3.3.5 Conformal-modular fixed-point microphysics

After Section 2.3, the remaining gap is narrower. On the local finite-constraint MaxEnt branch of the third OPH axiom, the logarithm of the selected state is a quasi-local UV generator, and the refinement-stable branch lies inside one common finite-dimensional multiplier family. The unresolved question is which algebraic phase the refinement limit realizes. At each regulator stage the patch and cap algebras are type-I matrix algebras, but the finite regulator class is not refinement-closed, so the scaling-limit observer algebra may leave that class. OPH therefore proves a local-Gibbs/refinement-stable branch with propagation and endpoint control, but it does *not* prove that every such branch lands in the Bisognano–Wichmann geometric phase or that MaxEnt alone selects a vacuum or canonical cap state. The conditional boundary is **T1+T3**: they identify the BW/geometric branch boundary, while the internal extraction of a realized scaling-limit cap pair depends on the derived carried-collar schedule above the transported fixed-local-collar Markov/faithfulness datum and the later weak-$*$ extraction plus GNS gluing step.

Local finite-constraint MaxEnt branch. The constraint family $\mathcal{C}$ is generated by finitely many gauge-invariant local densities $\{O_a(x)\}$ of UV range $O(\ell_{\mathrm{UV}})$, with the same finite label set retained under refinement. This is not an extra postulate beyond the third OPH axiom; it is that axiom unpacked at regulator level.

Theorem 2.6 (Local constraints imply a local-Gibbs form). If the MaxEnt constraints are expectations of finitely many quasi-local operators $\{O_a\}$ with bounded support size at scale ℓ_{UV}, then the entropy maximizer is

$$\omega \propto \exp\left(-\sum_a \lambda_a O_a\right),$$

so the MaxEnt generator $H_{\mathrm{MaxEnt}} = -\log \omega$ is a UV-range quasi-local sum. This is exactly the local-Gibbs form used later.

Proof. Standard exponential-family result: maximum entropy subject to linear constraints $\langle O_a \rangle = c_a$ yields the Gibbs state with Lagrange multipliers λ_a. QED.

Derived propagation control on the same branch. Because H_{MaxEnt} is a finite-range or quasi-local sum on a finite type-I regulator net, standard Lieb–Robinson estimates [6] apply to the automorphism group it generates, or to any branch generator lying in the same bounded-support algebraic closure. Thus there is a finite propagation velocity v_{LR} and constants C, ξ such that for local observables A_X, B_Y,

$$\|[\tau_t(A_X), B_Y]\| \le C\,\|A_X\|\,\|B_Y\|\,\min(|X|, |Y|)\,e^{-(d(X,Y)-v_{\mathrm{LR}}|t|)/\xi}.$$

Because changing a bounded interval endpoint only adds or removes an $O(|\Delta v|)$ collar of local terms once the central endpoint piece is separated off, the same local branch also gives bounded-interval endpoint-Lipschitz matrix elements,

$$|\langle\psi, (K[I'] - K[I])\phi\rangle| \le C_{\psi,\phi,I_{\max}}\,|I'\Delta I|,$$

used later in the null-modular bridge. So quasi-local propagation and endpoint control are not external regularity selectors; they are the local-constraint MaxEnt branch written in dynamical form.

Refinement-stable multiplier branch. Because the same finite constraint family is preserved under coarse-graining, the regulator states lie in one common finite-dimensional multiplier family rather than in unrelated state spaces at different cutoffs. The "refinement-stable" language used later therefore means persistence along the stable or fixed branch of this multiplier family. This state-side notion is enough to compare realized states across cutoffs. It does *not* by itself upgrade fixed-cutoff edge labels to transportable refinement-persistent sectors, and it does not show that the directed colimit of such sectors is nontrivial. Accordingly, whenever the gauge derivation speaks of a refinement-stable directed colimit of sectors, the entire transportable directed-system/fiber clause is the explicit content of **T7**; the third OPH axiom supplies only the realized state branch along which one asks whether such sector transport persists.

Scaling limit and algebraic type. The regulator presentation gives a family of finite type-I algebras

$$\mathcal{A}_\ell(C) \cong \mathcal{B}(\mathcal{H}_{C,\ell}).$$

A scaling limit of this family need not remain type I. In the continuum-QFT case of interest one expects the local limit algebra $\mathcal{A}_\infty(C)$ to be non-type-I, typically type III. Accordingly the fundamental modular datum in the limit is the automorphism group of the pair $(\mathcal{A}_\infty(C), \omega_\infty^C)$, not a density matrix inside $\mathcal{A}_\infty(C)$.

Proposition 2.6 (Geometric modular action on caps on the T1-selected BW branch). On the **T1**+**T3** branch, for any scaling-limit cap pair $(\mathcal{A}_\infty(C), \omega_\infty^C)$ furnished on the realized refinement-stable branch, let $\alpha_{\lambda_C(s)}$ be the automorphism induced by the standard cap-preserving conformal subgroup $\lambda_C(s)$. Then

$$\sigma_t^{\omega_\infty^C} = \alpha_{\lambda_C(2\pi t)}.$$

If the limit cap algebra happens to remain type I, this may be written as $K_C = 2\pi B_C$. In the generic continuum case, the same statement is an outer modular action on a non-type-I algebra.

Proof. **T1** identifies the scaling-limit cap modular group with the standard geometric cap subgroup up to normalization. The modular KMS condition fixes the normalization to 2π. If the limit algebra is type I, the automorphism statement may be represented by a modular Hamiltonian K_C; otherwise the automorphism statement is the full content. QED.

Alternative derivation via net regularity. The same modular-covariance property can also be read off from a scaling-limit support map when the net satisfies the outer-regularity / minimal-support condition used later. This does *not* choose the BW branch; it clarifies how the geometric labeling of the limit net is read once that branch is present.

(NR) Outer regularity / minimal support. For any operator O, the intersection of all connected regions P with $O \in \mathcal{A}(P)$ is again a connected region, denoted $\mathrm{supp}(O)$.

Proposition 2.7 (Modular covariance from net regularity). Under (NR), define for any region $R \subset C$

$$f_t^C(R) := \bigcup_{O \in \mathcal{A}(R)} \mathrm{supp}\Big(\sigma_t^{\omega,C}(O)\Big).$$

Then $\sigma_t^{\omega,C}(\mathcal{A}(R)) = \mathcal{A}(f_t^C(R))$, which is exactly the desired modular-covariance property.

Proof. Since $\sigma_t^{\omega,C}$ is an automorphism of $\mathcal{A}(C)$, and (NR) allows one to read support from the net labeling, the map $R \mapsto f_t^C(R)$ is well-defined and consistent. QED.

Null-surface modular structure. On the same **T1+T3** branch, the null-surface modular machinery narrows as follows:

- **Fixed-cutoff null-strip bridge.** The null-strip package proves transferred cut-center data, theorem-local inherited left/right strip structure, exact-or-controlled four-term strip additivity on one inherited strip model, renormalized endpoint control up to the weak tail generator, and the derived half-sided modular pair whose Borchers–Wiesbrock consequence is an explicit positive null-translation generator on its Stone domain.
- **Derived half-sided modular inclusion.** Nested null half-line algebras satisfy half-sided modular inclusion on the OPH geometric scaling branch by Corollary 5.2e; Borchers–Wiesbrock then identifies the positive null-translation generator on its Stone domain together with the affine half-line modular relation [10].
- **Weak continuity and finite variation.** The bounded-interval and half-line endpoint control follow from the local MaxEnt branch and define the weak tail generator at fixed cutoff. The **T1+T3** scaling branch supplies the continuum null-generator setting in which that weak-tail data can be matched to the geometric null modular action of the relativistic phase.

Constraint set specification. On the local finite-constraint branch, the "correct fixed-cap constraint set" becomes explicit: constraints are the local conserved charges of the symmetries used in the derivation:

1. **Edge/cap label constraints:** fix the distribution of collar-sector labels, equivalently $\langle L_C \rangle$ for each cap size, giving the area term.
2. **Gauge charges:** fix boundary flux or charge operators.
3. **Geometric (conformal) charges:** fix the expectation of the conformal Killing charges that preserve the cap, i.e. the generator B_C or its microscopic lattice approximation.

MaxEnt therefore selects the unique finite-stage invariant state compatible with those conserved charges. What it does *not* prove by itself is that the scaling-limit state is the vacuum or canonical cap state. The stronger statement is exactly **T1+T3**: they identify the BW/geometric branch boundary, while the internal extraction of a realized scaling-limit cap pair depends on the vanishing carried-collar schedule and the later weak-$*$ extraction plus GNS gluing step. When the manuscript uses the phrase "BW/canonical cap phase," it means only this **T1**-selected branch after such a cap pair has been emitted; if the limit algebra is non-type-I, the geometric modular action on that branch is outer.

QNEC internalization. QNEC has rigorous QFT proofs in broad settings [14]. On the **T1+T3** branch, the Recoverable Generalized Entropy axiom can be supported internally by:

- EC + MaxEnt derive $S_{\mathrm{gen}} = S_{\mathrm{bulk}} + \langle L_C \rangle$ (Section 5.4).
- Once the downstream null-stress bridge and Einstein branch are in place, focusing becomes a semiclassical consequence in the same scaling regime.

Summary. The CMFP package therefore separates the dependency structure into two layers:

- **Internal to OPH:** the local-Gibbs form, quasi-local propagation, bounded-interval endpoint control, and the refinement-stable notion of sector persistence, all from the local finite-constraint MaxEnt branch.
- **Theorem-external branch statement:** that the scaling limit lands in the Bisognano–Wichmann geometric modular class. On that branch the limit algebra may be non-type-I, the modular action may be outer, and the 2π normalization and later null half-sided-inclusion bridge follow.

The internal extension route is a two-step chain, and the pressure point is the first object. First, one needs a canonical scaling-limit observer cap-pair realization from transported marginals. The theorem contract visible in the manuscript is that the realized transported cap-local system packages the cap-local test family, the projectively compatible transported marginal family, and the asymptotic transport-equivalence certificate. The remaining emitted witness is a derived vanishing carried-collar schedule on fixed local collars, explicitly decomposed into the constructive-recovery remainder and the faithful modular-defect term, namely convergence of the carried quantity

$$r_{\mathrm{FR}}(\varepsilon_{n,m,\delta}) + 4\lambda_*^{-1}\, \delta^{\mathrm{M}}_{A_{n,m,\delta}:B_{n,m,\delta}:D_{n,m,\delta}}(\varepsilon_{n,m,\delta})$$

to zero after transport to each fixed local collar model; only then can local weak-$*$ extraction and GNS gluing emit the realized scaling-limit cap pair. Beneath that derived schedule, the sharper raw datum is the transported fixed-local-collar Markov/faithfulness package, and the single missing clause inside that lower datum is the eventual common floor on the finite modular-transport family feeding the faithful modular-defect term. No separate exact-Markov-reference faithfulness input is missing there: on one fixed collar model, the exact-Markov comparison marginals inherit the same eventual floor once the exact-Markov modulus tends to zero. The manuscript does not derive a refinement-uniform common floor on the finite modular-transport marginals; this single clause remains external to the emitted theorem chain. Second, one needs ordered cut-pair rigidity on that realized limit. That stage remains symbolic until the realized scaling-limit cap pair exists. The weak-$*$ extraction plus rigidity route is therefore an explicit internal extension route, not part of the proved theorem surface.

3.4 Overlap Consistency and Gluing

The constructive part of overlap consistency is the tree-gluing theorem below. The structural part is the origin of the gluing redundancy itself. On the ordinary or central-defect branch, gauge-as-gluing is the finite-regulator overlap redundancy of local chart presentations: once the overlap algebras are realized in finite-dimensional charts, overlap-consistent rechartings of a cut form a compact unitary transition system. The collar theorem of Section 2.3 should therefore be read with its boundary group G_Σ understood as shorthand for this derived compact boundary action, not as a model-specific add-on. On the genuinely noncentral branch, the same weak gluing data are encoded by a compact crossed-module change system, so the fixed-cutoff collar theorem upgrades to a higher-gauge statement rather than failing. The fixed-cutoff topological package is therefore closed on all three branches. A second structural point is UV underdetermination. If each patch Hilbert space is tensored with an inert finite ancillary factor and the observable patch algebras are embedded as $\mathcal{A}(P) \otimes \mathbf{1}$, then the physical observables, collar conditional mutual information, carried Markov errors, and quotient normal forms are unchanged. So OPH fixes the physical UV branch only modulo gauge or implementation hiding together with such ancillary stabilization, not a unique microscopic presentation.

3.4.1 Constructive gluing on tree covers

Theorem 3.1 (tree gluing). Let a rooted tree of patches satisfy a tree-ordered overlap structure and a tripartite factorization (A_k, B_k, C_k) at step k. If a target state ρ^* obeys $I(A_k : C_k \mid B_k) \leq \varepsilon_k$, then there exist recovery maps $\mathcal{R}_k$ such that

$$\|\rho^*_{A_k B_k C_k} - (\mathrm{id}_{A_k} \otimes \mathcal{R}_k)(\rho^*_{A_k B_k})\|_1 \leq \delta_k,$$

with

$$\delta_k = 2\sqrt{\ln 2\, \varepsilon_k}.$$

The iteratively glued state $\hat{\rho}$ satisfies

$$\|\hat{\rho} - \rho^*\|_1 \leq \min\left(2, \sum_{k=2}^{n} \delta_k\right).$$

Proof. Induct on k. The recovery error contracts under CPTP maps, so the errors add. QED.

3.4.2 Gauge-as-gluing and loops

At finite regulator scale, the fixed-cutoff gauge-as-gluing package identifies the overlap-consistency redundancy of local chart presentations. Choose finite-dimensional local presentations of the overlap algebras on a connected cut Σ. Because the overlap algebras are matrix algebras, any overlap-consistent change of chart is inner and is implemented by a unitary on the cut Hilbert space. Fixing a reference chart, the compact closure of the subgroup generated by all such recharting unitaries is the boundary gluing group K_Σ; before fixing the reference chart, the same data form a compact unitary groupoid.

Proposition 3.2a (Derived gauge-as-gluing at finite regulator). Let a finite regulator chart be chosen for the patches meeting along a connected interface Σ. Then overlap consistency determines a compact unitary transition system on the cut data. On the ordinary or central-defect branch, this transition system reduces to a compact boundary group K_Σ, and when the triple-overlap defect is central its projective composition law lifts to a genuine action of a compact central extension $\widehat{K_\Sigma}$. Gauge is therefore the overlap redundancy itself, not an extra primitive.

Proof. On a finite-dimensional matrix algebra every *-automorphism is inner, so each overlap-consistent recharting is conjugation by a unitary on the cut Hilbert space. The subgroup generated by those unitaries has compact closure inside a finite-dimensional unitary group. If triple-overlap defects are central, the resulting projective composition law lifts to a central extension. QED.

Lemma 3.2b (trees vs loops). If the patch adjacency graph is a tree, one can choose local charts h_i so that the overlap labels satisfy $g_{ij} = h_i^{-1} h_j$ on all edges. If loops exist, the loop holonomy

$$H(\gamma) = g_{i_1 i_2} g_{i_2 i_3} \cdots g_{i_n i_1}$$

is invariant under local frame changes. Nontrivial holonomy is the obstruction to global trivialization. QED.

Corollary 3.2c (Collar consequence on the ordinary or central-defect branch). Let $B_\delta = B_L \cup B_R$ be a collar around a cap boundary Σ, and set $\widehat{K}_\Sigma = K_\Sigma$ on the ordinary branch. Then the EC theorem of Section 2.3 has the derived interpretation

$$H_{B_\delta} = (\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{\widehat{K}_\Sigma} \cong \bigoplus_\alpha (H_{b_L^\alpha} \otimes H_{b_R^\alpha}),$$

with

$$Z(A(B_\delta)) = \bigoplus_\alpha \mathbb{C}\,\mathbf{1}_\alpha.$$

The right half-collar carries the contragredient action because it sees the inverse transport across the same cut. Exact Markov normal forms used later require the additional state hypothesis $I_\omega(A_\delta : D_\delta \mid B_\delta) = 0$ or the explicitly stated idealized recoverability limit; the block decomposition itself is kinematic and follows from the derived boundary action.

Proof sketch. Decompose the left boundary data into irreps $(V_\alpha \otimes H_{b_L^\alpha})$ of $\widehat{K}_\Sigma$ and the right boundary data into the dual modules $(V_\beta^* \otimes H_{b_R^\beta})$. Then Schur's lemma leaves a singlet only when $\alpha = \beta$, producing the displayed direct sum. QED.

3.4.3 Loop obstruction class (central defect)

On the central-defect subbranch, define central defects z_{ijk} by

$$\varphi_{ij}\varphi_{jk}\varphi_{ki} = \mathrm{Ad}(z_{ijk}) \quad \text{on } \mathcal{A}_{ijk}.$$

Then $\{z_{ijk}\}$ is a Čech 2-cocycle, and its cohomology class $[z]$ is gauge invariant. Central defects do not obstruct the collar block decomposition: they only replace K_Σ by its central extension $\widehat{K}_\Sigma$. Ordinary loop-coherent gluing exists iff $[z] = 0$. (A full proof appears in Section 6.4 below, in the algebra-net language.)

3.4.4 Non-central obstruction (2-group cocycle)

When defects are not central, the natural coefficient data is a crossed module $(H \to G)$ with an action of G on H by conjugation. Here G is the reconstructed gauge group, and H is the unitary group acting on edge multiplicity spaces, with boundary map $\partial : H \to G$.

A crossed module is a homomorphism $\partial : H \to G$ together with an action of G on H such that

$$\partial(g \triangleright h) = g\,\partial(h)\,g^{-1}, \qquad \partial(h) \triangleright h' = hh'h^{-1}.$$

On a good cover $\{P_i\}$, a weakly coherent gluing is encoded by:

$$g_{ij} : P_{ij} \to G, \qquad h_{ijk} : P_{ijk} \to H,$$

obeying the 2-cocycle conditions

$$g_{ij}g_{jk} = \partial(h_{ijk})g_{ik},$$

and on quadruple overlaps,

$$h_{jkl}h_{ijl} = (g_{ij} \triangleright h_{ikl})h_{ijk}.$$

Gauge changes act by 1- and 2-cochains in the standard way for crossed-module cohomology.

Theorem 3.4 (non-central obstruction). Loop-coherent gluing exists iff the 2-cocycle (g_{ij}, h_{ijk}) is equivalent to the trivial cocycle in nonabelian Čech H^2 with values in the crossed module $(H \to G)$.

Proof sketch. Strict gluing corresponds to $h_{ijk} = 1$ and $g_{ij}g_{jk} = g_{ik}$. Gauge changes are exactly the crossed-module coboundaries, so strictification exists iff the 2-class is trivial. QED.

The central-defect case is the abelian truncation with H central and trivial action, which reduces to Section 3.3. A genuinely noncentral class is the point where the fixed-cutoff collar theorem upgrades from the ordinary-group package to the higher-gauge one below.

Corollary 3.4a (Fixed-cutoff higher-gauge EC and transportability). For a finite regulator chart on a connected cut Σ, the genuinely noncentral branch admits a compact crossed-module change system

$$\mathcal{T}_\Sigma = C^1(N_\Sigma, H_\Sigma) \rtimes C^0(N_\Sigma, G_\Sigma).$$

The physical higher-gauge collar is

$$\mathcal{H}_B^{2g} = (\widetilde{\mathcal{H}}_{B_L} \otimes \widetilde{\mathcal{H}}_{B_R})^{\mathcal{T}_\Sigma} \cong \bigoplus_\lambda (\mathcal{H}_{b_L^\lambda} \otimes \mathcal{H}_{b_R^\lambda}),$$

with

$$Z(\mathcal{A}_{2g}(B)) = \bigoplus_\lambda \mathbb{C}\,\mathbf{1}_\lambda,$$

and the defect class

$$q_\Sigma = [(g, h)] \in \check{H}^2(N_\Sigma, H_\Sigma \to G_\Sigma)$$

is invariant under local rechartings, classifies fixed-cutoff genuinely noncentral sectors, and vanishes iff the defect is removable.

Proof sketch. Pair-overlap rechartings are inner, while triple-overlap associators strictify to compact crossed-module data. Finite-dimensional unitary $\mathcal{T}_\Sigma$-modules decompose semisimply, and Schur matching leaves only diagonal left/right blocks. The transport statement is the crossed-module Čech analogue of the central-defect case. QED.

3.5 Modular Flow and Lorentz Kinematics

3.5.1 Modular additivity in the Markov collar limit

Consider a collar tripartition $A : B : D$ around a cap boundary, with the EC decomposition of Section 2.3 understood. Define, for a faithful reference state ω,

$$\Delta K(\omega) := K_{ABD}(\omega) - K_{AB}(\omega) - K_{BD}(\omega) + K_B(\omega).$$

Exact modular additivity is not a consequence of EC alone. It is available only on the exact Markov set, or along a controlled fixed-cutoff family that approaches that set on one fixed collar model. Three quantities must therefore be kept separate:

1. the raw collar conditional mutual information $I(A : D \mid B)$;

2. the constructive Fawzi–Renner comparison error

$$r_{\mathrm{FR}}(\varepsilon) := 2\sqrt{1 - e^{-\varepsilon}} \leq 2\sqrt{\varepsilon};$$

3. the fixed-collar exact-Markov replacement modulus

$$\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) := \sup\left\{\inf_{\sigma\in\mathfrak{M}_{A:B:D}} \|\rho-\sigma\|_1 : I(A:D\mid B)_\rho \le \varepsilon\right\},$$

where

$$\mathfrak{M}_{A:B:D} := \{\sigma_{ABD} : I(A:D\mid B)_\sigma = 0\}.$$

Lemma 4.1a (Exact Markov implies exact additivity up to a central term). If $I(A:D\mid B)_\omega = 0$, then the EC decomposition of Section 2.3 puts the state in HJPW block form

$$\omega_{ABD} = \bigoplus_\alpha p_\alpha\, \omega^{(\alpha)}_{Ab^\alpha_L} \otimes \omega^{(\alpha)}_{b^\alpha_R D},$$

and $\Delta K(\omega)$ is central. On the canonical HJPW block model one may take

$$\Delta K(\omega) = 0.$$

Equivalently, there exists a central operator $K_{\partial,ABD}(\omega) \in Z(\mathcal{A}(B))$ such that

$$K_{ABD}(\omega) = K_{AB}(\omega) + K_{BD}(\omega) - K_B(\omega) + K_{\partial,ABD}(\omega).$$

Proof. On each HJPW block the modular Hamiltonians of ABD, AB, BD, and B are the logarithms of tensor-product states with the same classical block weight p_α. The blockwise logarithms therefore cancel exactly in the combination $K_{ABD} - K_{AB} - K_{BD} + K_B$. If one keeps the blockwise endpoint-label bookkeeping explicit rather than absorbing it into the canonical block identification, the remainder is a direct sum of block constants, hence central. QED.

Proposition 4.1b (Controlled exact-Markov replacement on a fixed collar). Because the state space is compact and conditional mutual information is continuous in finite dimension,

$$\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) \longrightarrow 0 \qquad (\varepsilon \downarrow 0).$$

Thus small collar CMI converges to the exact HJPW normal form only in this controlled fixed-cutoff sense; the manuscript does not claim a dimension-free one-shot trace-norm bound from small $I(A:D\mid B)$ directly to an exact Markov state.

Corollary 4.1c (Carried collar defect operator). Let ω_ε satisfy $I(A:D\mid B)_{\omega_\varepsilon} \le \varepsilon$, and choose $\sigma_\varepsilon \in \mathfrak{M}_{A:B:D}$ with

$$\|\omega_\varepsilon - \sigma_\varepsilon\|_1 \le \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

Define the carried defect operator

$$\mathfrak{D}_{A:B:D}(\omega_\varepsilon,\sigma_\varepsilon) := \Delta K(\omega_\varepsilon) - \Delta K(\sigma_\varepsilon).$$

Then

$$K_{ABD}(\omega_\varepsilon) = K_{AB}(\omega_\varepsilon) + K_{BD}(\omega_\varepsilon) - K_B(\omega_\varepsilon) + K_{\partial,ABD}(\sigma_\varepsilon) + \mathfrak{D}_{A:B:D}(\omega_\varepsilon,\sigma_\varepsilon),$$

where $K_{\partial,ABD}(\sigma_\varepsilon) = \Delta K(\sigma_\varepsilon)$ is central. For every bounded observable X supported on $A\cup B$,

$$|\mathrm{Tr}[X(\omega_\varepsilon - \sigma_\varepsilon)]| \le \|X\|_\infty\, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

If the relevant marginals of ω_ε and σ_ε are uniformly faithful with lower spectral bound $\lambda_* > 0$, then

$$\|\mathfrak{D}_{A:B:D}(\omega_\varepsilon,\sigma_\varepsilon)\|_\infty \le 4\lambda_*^{-1}\, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

Independently, Fawzi–Renner recovery supplies a constructive comparison state $\omega_\varepsilon^{\mathrm{rec}}$ with

$$\|\omega_\varepsilon - \omega_\varepsilon^{\mathrm{rec}}\|_1 \le r_{\mathrm{FR}}(\varepsilon), \qquad r_{\mathrm{FR}}(\varepsilon) := 2\sqrt{1-e^{-\varepsilon}} \le 2\sqrt{\varepsilon}.$$

This $O(\varepsilon^{1/2})$ term is the finite-stage observable error. The modulus $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon)$ is instead the fixed-collar quantity that justifies replacing the physical state by an exact Markov normal form in the later geometric modular arguments.

Accordingly, the manuscript's later exact collar formulas take one of two forms:

1. literal exactness when the reference state is exact Markov on the relevant collar; or
2. a controlled collar family for which $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon_\delta) \to 0$, with the constructive Fawzi–Renner remainder and the carried operator defect $\mathfrak{D}_{A:B:D}$ kept explicit at finite stage.

3.5.2 Theorem: BW_{S^2} on the OPH geometric branch

The collar analysis proves a fixed-cutoff statement on the finite type-I regulator net: the reduced cap state has a literal density matrix, its modular Hamiltonian exists, and its nonadditive part is confined to a shrinking collar up to carried errors. The Lorentz claim is therefore not a fixed-cutoff matrix-algebra identity. Its target is the refinement-limit cap pair $(\mathcal{A}_\infty(C), \omega_\infty^C)$ singled out by the controlled scaling clause **T3**. Nothing in the local MaxEnt branch by itself proves that the scaling-limit cap algebra remains type I, and nothing in the fixed-cutoff MaxEnt/collar package by itself proves that the realized scaling branch lands in the Bisognano–Wichmann geometric modular class. The role of **T1** is exactly to identify that branch.

Fix a cap $C \subset S^2$ and a shrinking collar family $(A_\delta, B_\delta, D_\delta)$ around ∂C. Write

$$\varepsilon_\delta := I(A_\delta : D_\delta \mid B_\delta)_\omega, \qquad r_{\mathrm{FR}}(\varepsilon_\delta) := 2\sqrt{1-e^{-\varepsilon_\delta}} \le 2\sqrt{\varepsilon_\delta},$$

and, on each fixed faithful collar model,

$$\eta_\delta^{\mathrm{M}} := 4\lambda_*^{-1}\, \delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta),$$

where $\lambda_* > 0$ is the lower spectral bound used to compare modular Hamiltonians. All exact collar formulas in this subsection are therefore to be read either literally at exact Markovity or along a controlled collar family satisfying

$$\delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta) \to 0 \qquad (\delta \downarrow 0),$$

with $r_{\mathrm{FR}}(\varepsilon_\delta)$ and η_δ^{M} carried explicitly.

For each cap C, let $\lambda_C(s) \subset \mathrm{Conf}^+(S^2)$ denote the standard cap-preserving conformal one-parameter subgroup, normalized so that the null blow-up near a smooth cut acts by $v \mapsto e^{-s}v$, and let $\alpha_{\lambda_C(s)}$ denote the induced automorphism of the scaling-limit cap net.

Theorem 4.2 (BW_{S^2} on the OPH Bisognano–Wichmann branch). Let $C \subset S^2$ be a round cap. Assume the OPH axioms, the derived fixed-cutoff regulator/collar package established earlier, the controlled scaling-limit scope clause **T3**, and the branch condition **T1** that the realized refinement-stable scaling branch for caps and their conformal images lies in the Bisognano–Wichmann geometric modular class. Then:

(i) at each finite regulator stage, the cap algebra is type I, the reduced cap state has a density matrix $\rho_C^{(\delta)}$, and its modular Hamiltonian

$$K_C^{(\delta)} := -\log \rho_C^{(\delta)}$$

has nonadditive part confined to the shrinking collar up to the carried errors $r_{\mathrm{FR}}(\varepsilon_\delta)$ and η_δ^{M};

(ii) for every realized scaling-limit cap pair $(\mathcal{A}_\infty(C), \omega_\infty^C)$ on that branch, the modular automorphism group is geometric:

$$\sigma_t^{\omega_\infty^C} = \alpha_{\lambda_C(2\pi t)};$$

(iii) no separate cap-isotropy/SO(2)-equivariance selector or Euclidean-regularity premise is used here: once **T1** has fixed the geometric branch and $\lambda_C(s)$ is normalized by the null blow-up $v \mapsto e^{-s}v$, the modular KMS normalization fixes the coefficient 2π. No claim is made that $\mathcal{A}_\infty(C)$ remains type I. If the scaling-limit cap algebra remains type I, one may write

$$K_C = 2\pi B_C.$$

In the generic continuum case, the automorphism identity in item (ii) is the correct theorem statement and the geometric modular action is generally outer rather than inner.

Consequently, for every controlled collar family with

$$\delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta) \to 0$$

and surviving faithfulness bound, replacing the finite-stage modular action by the geometric cap-dilation action incurs only the carried errors $r_{\mathrm{FR}}(\varepsilon_\delta)$ and η_δ^{M}, and these vanish in the refinement limit.

Proof. The fixed-cutoff collar theorems localize the modular defect to the shrinking collar and control the replacement of the physical collar state by an exact-Markov reference with the carried errors $r_{\mathrm{FR}}(\varepsilon_\delta)$ and η_δ^{M}. By **T3** the theorem is about the scaling-limit cap pair rather than the finite regulator matrices. By **T1**, on the selected branch the scaling-limit modular automorphism group is the standard cap-preserving conformal flow $\alpha_{\lambda_C(s)}$ up to normalization:

$$\sigma_t^{\omega_\infty^C} = \alpha_{\lambda_C(\kappa_C t)}.$$

Because **T1** identifies the standard cap-preserving subgroup, no separate cap-isotropy or Euclidean-regularity selector enters. The modular group is KMS at inverse temperature 1 by definition. The modular KMS condition together with the null-blow-up normalization $v \mapsto e^{-s}v$ fixes $\kappa_C = 2\pi$. If the limit algebra remains type I, one may represent the automorphism by a modular Hamiltonian and write $K_C = 2\pi B_C$; otherwise the automorphism identity is the full statement. QED.

BW boundary. The theorem proves geometric modular flow only on the **T1**-selected scaling branch. It does *not* prove that MaxEnt alone selects that branch, nor that the continuum state is a vacuum or canonical cap state. When the manuscript informally uses the phrase “BW/canonical cap phase,” it means only this selected scaling-limit branch, and in the generic non-type-I case its geometric modular action is outer.

The internal extension route is a two-object scaffold. First, one needs a realized scaling-limit cap-pair extraction from transported marginals. The theorem contract visible in the manuscript is that the realized transported cap-local system packages the cap-local test family, the projectively compatible transported marginal family, and the asymptotic transport-equivalence certificate. The remaining emitted witness is a derived vanishing carried-collar schedule on fixed local collars, explicitly decomposed into the constructive-recovery remainder and the faithful modular-defect term, namely the convergence to zero of

$$r_{\mathrm{FR}}(\varepsilon_{n,m,\delta}) + 4\lambda_*^{-1}\,\delta^{\mathrm{M}}_{A_{n,m,\delta}:B_{n,m,\delta}:D_{n,m,\delta}}(\varepsilon_{n,m,\delta})$$

after transport to each fixed local collar model; only then can local weak-$*$ extraction and GNS gluing emit the realized scaling-limit cap pair. Beneath that derived schedule, the sharper raw datum is the transported fixed-local-collar Markov/faithfulness package, and the single missing clause inside that lower datum is the eventual common floor on the finite modular-transport family feeding the faithful modular-defect term. No separate exact-Markov-reference faithfulness input is missing there: on one fixed collar model, the exact-Markov comparison marginals inherit the same eventual floor once the exact-Markov modulus tends to zero. The manuscript does not derive a refinement-uniform common floor on the finite modular-transport marginals; this single clause remains external to the emitted theorem chain. Second, one needs ordered cut-pair rigidity on that realized limit. That rigidity step remains symbolic until the realized scaling-limit cap pair exists.

3.5.3 Theorem: BW_{S^2} implies Lorentz kinematics

Theorem 4.3 (Lorentz kinematics on the screen). Under the hypotheses of Theorem 3.5.2,

$$\mathrm{Conf}^+(S^2) \cong \mathrm{PSL}(2,\mathbb{C}) \cong \mathrm{SO}^+(3,1).$$

The cap modular flows are therefore the standard one-parameter Lorentz boost/dilation subgroups in the celestial-sphere realization, and the induced local kinematic group is the connected Lorentz group.

Proof. Orientation-preserving conformal maps of S^2 are exactly the Möbius transformations, so

$$\mathrm{Conf}^+(S^2) \cong \mathrm{PSL}(2,\mathbb{C}) \cong \mathrm{SO}^+(3,1).$$

By Theorem 3.5.2, on the **T1** branch the realized scaling-limit cap modular automorphism groups are the standard cap-preserving conformal dilations with the 2π normalization fixed internally. Hence the local kinematics induced by modular flow is the connected Lorentz group. QED.

3.6 Gravity from Entanglement Equilibrium

3.6.1 Cap first law

Section 4.2 fixes the cap modular statement first at the automorphism level. On the **T1** branch, the scaling-limit cap modular flow is geometric with the standard 2π normalization,

$$\sigma_t^{\omega_\infty^C} = \alpha_{\lambda_C(2\pi t)}.$$

If the realized scaling-limit cap algebra happens to remain type I, one may choose a density matrix ρ_C^ω and modular Hamiltonian

$$K_C := -\log \rho_C^\omega,$$

and then for a perturbation $\rho(\varepsilon)$ with $\rho(0) = \omega$ the first-law identity reads

$$\delta S_C = \delta\langle K_C\rangle = 2\pi\, \delta\langle B_C\rangle.$$

In the generic continuum/non-type-I case, Section 4.2 does not supply an inner operator identity $K_C = 2\pi B_C$; the theorem surface there is only the geometric modular automorphism statement, so the cap first-law formula is conditional on the special type-I realization.

3.6.2 Null-surface modular bridge

This subsection extends the fixed-cutoff weak-tail-generator boundary to the derived positive null-translation stage. At fixed cutoff one can transfer cut-center data to regulated null strips, obtain exact or controlled strip additivity on one inherited strip model, and derive a weak tail generator for the renormalized half-line family. On the OPH geometric scaling branch of Theorem 4.2, the null half-line blow-up net then inherits geometric dilation and therefore half-sided modular inclusion, so Borchers–Wiesbrock supplies an explicit positive null-translation generator on its Stone domain. The downstream density upgrade and the relativistic null-stress identification are outside this subsection. Accordingly the c.12-local bridge consists of:

- null-cut center transfer and the inherited split condition;
- exact or controlled four-term strip additivity on one fixed inherited strip model;
- endpoint-Lipschitz control for the renormalized half-line family and the resulting weak tail generator;
- derived half-sided modular inclusion on the null half-line blow-up net, and the resulting Borchers positive translation generator.

The later density formula and T_{kk} bridge are kept below only as explicit downstream templates rather than as fixed-cutoff theorems proved in this subsection. Lemma 5.2f makes the positive null-translation generator itself explicit: it records the Borchers unitary group, positivity and self-adjointness on the Stone domain, the corrected affine commutator relation, and the half-line modular-Hamiltonian identity $K_a = K_0 - 2\pi a P_\Omega$. Bounded-interval formulas remain downstream.

Proposition 5.2a (Null-cut center transfer and inherited split). Fix a regulated null tripartition

$$I_- = (v_1, v_2), \qquad J = (v_2, v_3), \qquad I_+ = (v_3, v_4),$$

with cuts $\Gamma_- := \{v = v_2\}$ and $\Gamma_+ := \{v = v_3\}$. Under the fixed-cutoff realized presentation, assume that the two null cuts inherit the ordinary or central-defect boundary-redundancy data used in the spatial collar branch, so that in a compatible type-I regulator presentation one has

$$\tilde{\mathcal{H}}_{I_-} \cong \bigoplus_{\alpha_-} W_{\alpha_-} \otimes \mathcal{H}_{i_-^{\alpha_-}},$$

$$\tilde{\mathcal{H}}_J \cong \bigoplus_{\alpha_-,\alpha_+} W^*_{\alpha_-} \otimes \mathcal{H}_{j^{\alpha_-,\alpha_+}} \otimes W_{\alpha_+},$$

$$\tilde{\mathcal{H}}_{I_+} \cong \bigoplus_{\alpha_+} W^*_{\alpha_+} \otimes \mathcal{H}_{i_+^{\alpha_+}},$$

where opposite sides of each cut carry inverse transport. If the strip algebra is the commutant of the transported cut actions, then

$$\mathcal{A}(J) \cong \bigoplus_{\alpha_-,\alpha_+} \mathcal{B}(\mathcal{H}_{j^{\alpha_-,\alpha_+}}), \qquad Z(\mathcal{A}(J)) = \bigoplus_{\alpha_-,\alpha_+} \mathbb{C}\, P_{\alpha_-,\alpha_+}.$$

If, in addition, each multiplicity space factors as

$$\mathcal{H}_{j^{\alpha_-,\alpha_+}} \cong \mathcal{H}_{j_L^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{j_R^{\alpha_-,\alpha_+}},$$

with $\mathcal{A}(I_- \cup J)$ acting blockwise only on $\mathcal{H}_{i_-^{\alpha_-}} \otimes \mathcal{H}_{j_L^{\alpha_-,\alpha_+}}$ and $\mathcal{A}(J \cup I_+)$ acting blockwise only on $\mathcal{H}_{j_R^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{i_+^{\alpha_+}}$, then

$$\mathcal{A}(J) \cong \bigoplus_{\alpha_-,\alpha_+} \mathcal{B}(\mathcal{H}_{j_L^{\alpha_-,\alpha_+}}) \otimes \mathcal{B}(\mathcal{H}_{j_R^{\alpha_-,\alpha_+}}),$$

and

$$\mathcal{H}_{I_-\cup J\cup I_+} \cong \bigoplus_{\alpha_-,\alpha_+} \mathcal{H}_{i_-^{\alpha_-}} \otimes \mathcal{H}_{j_L^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{j_R^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{i_+^{\alpha_+}}.$$

Proof. Complete reducibility gives the displayed decompositions. Commuting with the two transported cut actions forces strip operators to preserve the sector pair (α_-, α_+) and act only on the multiplicity space; Schur's lemma kills the off-diagonal intertwiners and leaves the direct-sum algebra above. Taking invariants across both cuts in the glued tripartition again uses Schur's lemma and leaves exactly the matching sector pairs. The left/right split of the multiplicity spaces is additional structure; it is not forced by the center transfer alone. QED.

Corollary 5.2b (Exact or controlled four-term null modular relation on an inherited strip model). On one fixed finite-dimensional strip model satisfying Proposition 5.2a, define

$$\mathfrak{M}_{I_-:J:I_+} := \{\sigma : \ I(I_- : I_+ \mid J)_\sigma = 0\},$$

and

$$\delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon) := \sup\left\{ \inf_{\sigma\in\mathfrak{M}_{I_-:J:I_+}} \|\rho-\sigma\|_1 : I(I_- : I_+ \mid J)_\rho \le \varepsilon \right\}.$$

For any strip state η, let

$$\Delta K_J(\eta) := K_{I_-\cup J\cup I_+}(\eta) - K_{I_-\cup J}(\eta) - K_{J\cup I_+}(\eta) + K_J(\eta).$$

If the strip reference state ω is exact Markov, then $\Delta K_J(\omega)$ is central, and on the canonical inherited HJPW model one may take

$$\Delta K_J(\omega) = 0.$$

Equivalently, there exists a central operator $K_{\partial,J}(\omega) \in Z(\mathcal{A}(J))$ such that

$$K_{I_-\cup J\cup I_+}(\omega) = K_{I_-\cup J}(\omega) + K_{J\cup I_+}(\omega) - K_J(\omega) + K_{\partial,J}(\omega).$$

If instead $I(I_- : I_+ \mid J)_\omega \le \varepsilon$, choose $\widetilde{\omega}_J \in \mathfrak{M}_{I_-:J:I_+}$ with

$$\|\omega - \widetilde{\omega}_J\|_1 \le \delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon).$$

Define

$$\mathfrak{D}_J(\omega, \widetilde{\omega}_J) := \Delta K_J(\omega) - \Delta K_J(\widetilde{\omega}_J).$$

Then

$$K_{I_-\cup J\cup I_+}(\omega) = K_{I_-\cup J}(\omega) + K_{J\cup I_+}(\omega) - K_J(\omega) + K_{\partial,J}(\widetilde{\omega}_J) + \mathfrak{D}_J(\omega, \widetilde{\omega}_J),$$

with $K_{\partial,J}(\widetilde{\omega}_J) = \Delta K_J(\widetilde{\omega}_J) \in Z(\mathcal{A}(J))$. Every bounded observable on $I_- \cup J$ or $J \cup I_+$ then differs from the exact-Markov reference by at most

$$\|O\|_\infty \, \delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon).$$

If the relevant strip marginals are uniformly faithful with lower spectral bound $\lambda_* > 0$, then

$$\|\mathfrak{D}_J(\omega, \widetilde{\omega}_J)\|_\infty \le 4\lambda_*^{-1} \, \delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon).$$

Independently, Fawzi–Renner gives a recovered comparison state with trace-norm error

$$r_{\mathrm{FR}}(\varepsilon) = 2\sqrt{1-e^{-\varepsilon}} \le 2\sqrt{\varepsilon}.$$

Thus the exact four-term strip relation is available only at exact Markovity or in a controlled strip family on one fixed inherited strip model with $\delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon_J) \to 0$, while $\mathfrak{D}_J$ is carried explicitly at finite stage.

Proof. On the inherited split of Proposition 5.2a, the exact-Markov case is the same HJPW block calculation as in the spatial collar theorem: the four logarithms cancel blockwise on the canonical model, and any residual bookkeeping term is central. The controlled replacement is the same fixed-model compactness argument used for ordinary collars, after relabeling $A, B, D \mapsto I_-, J, I_+$; the operator-norm bound on $\mathfrak{D}_J$ is the corresponding relabeling of the modular-transport estimate. QED.

Definition (Renormalized null modular functional). For a null interval I on generator Ω, let $K_\partial(I, \Omega)$ denote the central endpoint-label term singled out by Corollary 5.2b on the inherited strip model, or by its controlled exact-Markov replacement when that model is used as reference. Define

$$\widetilde{K}[I, \Omega] := K[I, \Omega] - K_\partial(I, \Omega), \qquad \widetilde{K}_a(\Omega) := \widetilde{K}[(a, \infty), \Omega].$$

Proposition 5.2c (Endpoint-Lipschitz null modular families and weak tail generator). For renormalized modular Hamiltonians on one null generator,

$$\widetilde{K}[I, \Omega] := K[I, \Omega] - K_\partial(I, \Omega), \qquad \widetilde{K}_a(\Omega) := \widetilde{K}[(a, \infty), \Omega],$$

the branch-internal endpoint-control coming from the local finite-constraint MaxEnt branch gives:

$$\left|\langle \psi, (\widetilde{K}[(a', b'), \Omega] - \widetilde{K}[(a, b), \Omega])\phi \rangle\right| \le C_{\psi,\phi,\Omega}\big(|a' - a| + |b' - b|\big)$$

for bounded intervals in a compact endpoint window, and

$$\left|\langle \psi, (\widetilde{K}_{a'}(\Omega) - \widetilde{K}_a(\Omega))\phi \rangle\right| \le C_{\psi,\phi,\Omega}|a' - a|$$

for half-lines. Hence

$$f_{\psi,\phi}(a) := \langle \psi, \widetilde{K}_a(\Omega)\phi \rangle$$

is locally Lipschitz and therefore absolutely continuous, and its distributional derivative defines the weak tail generator

$$\langle \psi, q(a, \Omega)\phi \rangle := -\frac{1}{2\pi}\, \partial_a f_{\psi,\phi}(a).$$

For $a < b$,

$$\langle \psi, (\widetilde{K}_b(\Omega) - \widetilde{K}_a(\Omega))\phi \rangle = -2\pi \int_a^b \langle \psi, q(v, \Omega)\phi \rangle \, dv.$$

At this stage $q(a, \Omega)$ is a weak tail generator, not yet a local operator-valued density; its identification with the positive self-adjoint Borchers generator occurs only after Corollary 5.2e and Lemma 5.2f.

Proof. The derived endpoint-control estimate applies to the renormalized family after removal of the central endpoint term. For bounded intervals, the symmetric-difference length is bounded by $|a' - a| + |b' - b|$; for half-lines it is exactly $|a' - a|$. Therefore the matrix-element functions are Lipschitz in the endpoints. A Lipschitz function on a compact interval is absolutely continuous, so $f_{\psi,\phi}(a)$ has an L^∞_{loc} derivative and obeys the fundamental theorem of calculus. Defining q as the rescaled negative derivative gives the displayed integral relation. QED.

Downstream c.12–c.13 boundary. The fixed-cutoff bridge established above is the end of this subsection's proved content.

Lemma 5.2d (Downstream density-upgrade template). If, in addition, the weak tail generator $q(a,\Omega)$ is weakly differentiable in a and both $f_{\psi,\phi}(a)$ and $q_{\psi,\phi}(a)$ vanish at $+\infty$, then one may define a density

$$\langle\psi, p(a,\Omega)\phi\rangle := -\partial_a\langle\psi, q(a,\Omega)\phi\rangle = \frac{1}{2\pi}\,\partial_a^2\langle\psi, \widetilde{K}_a(\Omega)\phi\rangle$$

and obtain the half-line formula

$$\langle\psi, \widetilde{K}_a(\Omega)\phi\rangle = 2\pi\int_a^\infty (v-a)\,\langle\psi, p(v,\Omega)\phi\rangle\,dv.$$

This is a downstream density-upgrade step and is not proved from the fixed-cutoff strip arguments above.

Null half-line blow-up net. Fix a smooth cut point and choose affine coordinate v on generator Ω so that the cut sits at $v=0$. For $a \geq 0$, write

$$H_a := (a,\infty), \qquad \mathcal{M}_a(\Omega) := \overline{\bigvee_{a<c<d<\infty} \mathcal{A}((c,d),\Omega)}.$$

By isotony of the null interval net,

$$a \leq b \quad\Longrightarrow\quad \mathcal{M}_b(\Omega) \subseteq \mathcal{M}_a(\Omega).$$

Corollary 5.2e (Derived half-sided modular inclusion on null half-lines). On the OPH geometric scaling branch of Theorem 4.2, the blow-up modular action near a smooth entangling cut acts on the null coordinate by

$$v \mapsto e^{-2\pi t}v.$$

Therefore, for every $a \geq 0$,

$$\sigma_t^\omega(\mathcal{M}_a(\Omega)) = \mathcal{M}_{e^{-2\pi t}a}(\Omega).$$

Hence, for every $a > 0$,

$$\sigma_t^\omega(\mathcal{M}_a(\Omega)) \subseteq \mathcal{M}_a(\Omega) \qquad (t \leq 0),$$

so the inclusion $\mathcal{M}_a(\Omega) \subset \mathcal{M}_0(\Omega)$ is half-sided modular. After the harmless convention change $t \mapsto -t$, this is the standard positive-time half-sided-inclusion form. The half-sided modular inclusion is therefore derived from OPH null-interval structure, isotony, and the scaling-limit geometric action rather than imported separately.

Proof. Blow up the cap modular flow of Theorem 4.2 near a smooth cut and restrict to the chosen null generator. In the tangent limit the cap-preserving flow becomes the null dilation $v \mapsto e^{-2\pi t}v$. For every bounded interval $(c,d) \subset H_a$, this sends $\mathcal{A}((c,d),\Omega)$ to $\mathcal{A}((e^{-2\pi t}c, e^{-2\pi t}d),\Omega)$, and taking the von Neumann closure of the interval net yields the displayed identity for $\mathcal{M}_a(\Omega)$. If $t \leq 0$, then $e^{-2\pi t}a \geq a$, so $H_{e^{-2\pi t}a} \subseteq H_a$, and isotony gives $\mathcal{M}_{e^{-2\pi t}a}(\Omega) \subseteq \mathcal{M}_a(\Omega)$. QED.

Lemma 5.2f (Positive null-translation generator). For the derived half-sided modular inclusion

$$\mathcal{M}_a(\Omega) \subset \mathcal{M}_0(\Omega) \qquad (a > 0),$$

let $\Delta_0(\Omega)$ be the modular operator of the standard pair $(\mathcal{M}_0(\Omega),\omega)$ and define $K_0(\Omega) := -\log\Delta_0(\Omega)$. Borchers–Wiesbrock then yields a unique strongly continuous one-parameter unitary group

$$U_\Omega(a) = e^{iaP_\Omega}, \qquad a \in \mathbb{R},$$

such that

$$U_\Omega(a)\omega = \omega, \qquad U_\Omega(a)\mathcal{M}_b(\Omega)U_\Omega(a)^* = \mathcal{M}_{a+b}(\Omega) \quad (a, b \geq 0),$$

and whose generator P_Ω is positive and self-adjoint on the Stone domain

$$D(P_\Omega) := \left\{ \psi \in \mathcal{H} : \lim_{a\to 0} \frac{U_\Omega(a)\psi - \psi}{ia} \text{ exists} \right\}.$$

Moreover,

$$\Delta_0(\Omega)^{it}U_\Omega(a)\Delta_0(\Omega)^{-it} = U_\Omega(e^{-2\pi t}a), \qquad \Delta_0(\Omega)^{it}P_\Omega\Delta_0(\Omega)^{-it} = e^{-2\pi t}P_\Omega,$$

and on the common invariant analytic core of $K_0(\Omega)$ and P_Ω,

$$[K_0(\Omega), P_\Omega] = -\, i\, 2\pi P_\Omega.$$

If $\Delta_a(\Omega)$ denotes the modular operator of $(\mathcal{M}_a(\Omega), \omega)$ and $K_a(\Omega) := -\log \Delta_a(\Omega)$, then

$$\Delta_a(\Omega) = U_\Omega(a)\Delta_0(\Omega)U_\Omega(a)^*, \qquad K_a(\Omega) = U_\Omega(a)K_0(\Omega)U_\Omega(a)^*,$$

so

$$K_a(\Omega) = K_0(\Omega) - 2\pi a\, P_\Omega$$

as a quadratic-form identity on $D(K_0(\Omega)) \cap D(P_\Omega)$. Equivalently,

$$\langle \psi, (K_b(\Omega) - K_a(\Omega))\phi \rangle = -2\pi(b-a)\langle \psi, P_\Omega \phi \rangle$$

for all $\psi, \phi \in D(K_0(\Omega)) \cap D(P_\Omega)$. In the canonical normalization of Corollary 5.2b, the weak endpoint derivative from Proposition 5.2c is therefore exactly the Borchers generator. Bounded-interval modular-Hamiltonian formulas remain downstream and require the separate interval-preserving branch recorded in Theorem 5.2g.

Proof. Corollary 5.2e gives the required half-sided modular inclusion. Borchers–Wiesbrock then yields the unique strongly continuous unitary group $U_\Omega(a)$, its positive self-adjoint generator P_Ω, and the affine covariance relation with the modular group; Stone's theorem gives the displayed domain formula. Differentiating

$$\Delta_0(\Omega)^{it}U_\Omega(a)\Delta_0(\Omega)^{-it} = U_\Omega(e^{-2\pi t}a)$$

first at $a = 0$ and then at $t = 0$ on the common analytic core gives

$$[K_0(\Omega), P_\Omega] = -\, i\, 2\pi P_\Omega.$$

Because $U_\Omega(a)\omega = \omega$ and $U_\Omega(a)\mathcal{M}_0(\Omega)U_\Omega(a)^* = \mathcal{M}_a(\Omega)$, uniqueness of modular data for the translated standard pair gives the displayed formulas for $\Delta_a(\Omega)$ and $K_a(\Omega)$. Differentiating $K_a(\Omega) = U_\Omega(a)K_0(\Omega)U_\Omega(a)^*$ in a yields $dK_a/da = -2\pi P_\Omega$ as a quadratic-form identity, and integrating from 0 to a gives the stated half-line modular-Hamiltonian relation. QED.

Theorem 5.2g (Downstream null-stress identification template). If one adds the density-upgrade hypotheses of Lemma 5.2d and the relativistic scaling-limit null-stress premise, then on the half-line blow-up net whose half-sided modular inclusion is derived in Corollary 5.2e and whose positive null-translation generator is identified in Lemma 5.2f, one expects an operator-valued distribution $T_{kk}(v, \Omega)$ with

$$\widetilde{K}[(a, \infty), \Omega] = 2\pi \int_a^\infty (v - a)\, T_{kk}(v, \Omega)\, dv + E^{(\eta)}_{(a,\infty),\Omega},$$

and, for bounded null intervals that inherit geometric modular action,

$$\widetilde{K}[I,\Omega] = 2\pi \int_I g_I(v)\, T_{kk}(v,\Omega)\, dv + E^{(\eta)}_{I,\Omega}.$$

Here, for each relevant half-line or bounded interval I, ε_I denotes the strip-CMI control parameter, δ_I^{M} the corresponding fixed-strip exact-Markov modulus from Corollary 5.2b, and

$$\eta_I := r_{\mathrm{FR}}(\varepsilon_I) + 4\lambda_*^{-1}\delta_I^{\mathrm{M}}(\varepsilon_I) + \eta_I^{\mathrm{EFT}},$$

so that

$$|\langle \psi, E^{(\eta)}_{I,\Omega}\phi\rangle| \le C_{I,\Omega,\psi,\phi}\, \eta_I$$

for domain vectors ψ, ϕ. The bounded-interval kernel $g_I(v)$ is affine-covariant and vanishes at finite interval endpoints. The full modular Hamiltonian is

$$K[I,\Omega] = \widetilde{K}[I,\Omega] + K_\partial(I,\Omega),$$

with K_∂ central and boundary-supported, and knowledge of T_{kk} for all null directions reconstructs a symmetric tensor T_{ab} up to a metric term. This subsection does not prove this theorem; it records the exact downstream burden beyond the derived half-sided-inclusion bridge.

Boundary of the conditional null bridge. The null bridge is conditional rather than automatic. Its explicit burden is:

- transferred cut-center data alone give the central sector decomposition of the null strip, not yet the left/right HJPW split;
- the extra decomposition-inheritance condition of Proposition 5.2a is exactly what upgrades those sectors to the same collar-type block structure used in the spatial branch;
- the fixed-cutoff null-strip bridge is exact for exact Markov strip states on that inherited decomposition, and otherwise is replaced by a controlled limit with the carried defect operator $\mathfrak{D}_J$;
- the local finite-constraint MaxEnt branch gives endpoint-Lipschitz control strong enough to define a weak tail generator for renormalized half-lines, and the OPH geometric scaling branch then derives half-sided modular inclusion on the half-line blow-up net;
- Borchers–Wiesbrock therefore supplies a genuine positive null-translation generator P_Ω inside the bridge itself, together with its Stone domain and the half-line modular-Hamiltonian relation $K_a = K_0 - 2\pi a P_\Omega$;
- the density upgrade and the final identification with a local continuum T_{kk} are the remaining explicit downstream steps.

This is the route from null-strip factorization to the later Jacobson branch, without identifying the null modular bridge itself with the downstream density-upgrade and null-stress steps.

3.6.3 Modular energy as stress-tensor charge (UV CFT)

If one assumes a UV CFT regime on sufficiently small caps, the modular Hamiltonian is explicitly local. This serves as an alternative EFT bridge to Section 5.2.

For a CFT vacuum on a ball, the modular Hamiltonian is local:

$$H_\zeta = \int_\Sigma T_{ab}\zeta^b d\Sigma^a,$$

where ζ is the conformal Killing field preserving the diamond. For a small diamond of size ℓ in $d = 4$,

$$\delta\langle H_\zeta\rangle = \frac{4\pi\ell^4}{15}\delta\langle T_{00}\rangle + O(\ell^5),$$

in the diamond rest frame.

3.6.4 Localized generalized entropy from Markov + MaxEnt

Using the collar decomposition and the collar double-scaling hypothesis, supplied either as a stated collar-limit premise or derived via Theorem 2.5 from the local-Gibbs form together with exponential mixing, the state takes the Markov normal form. MaxEnt selection maximizes entropy within each edge sector, producing

$$\rho_C = \bigoplus_\alpha p_\alpha \left(\rho^\alpha_{\text{bulk},C} \otimes \frac{\mathbf{1}^\alpha_{\text{edge}}}{d_\alpha}\right).$$

The entropy splits as

$$S(\rho_C) = H(p_\alpha) + \sum_\alpha p_\alpha S(\rho^\alpha_{\text{bulk},C}) + \sum_\alpha p_\alpha \log d_\alpha.$$

Convention: Throughout this paper, "log" denotes the natural logarithm (ln), so entropies are measured in **nats** (1 nat = 1/ln 2 ≈ 1.443 bits). This is standard in thermodynamics and QFT; the Bekenstein-Hawking formula S = A/4G uses nats. When clarity requires it, we write $\log_2$ explicitly for bits.

Define

$$S_{\text{bulk}}(C) := H(p_\alpha) + \sum_\alpha p_\alpha S(\rho^\alpha_{\text{bulk},C}),$$

and the central area operator

$$L_C := \sum_\alpha (\log d_\alpha) P_\alpha.$$

Then

$$S_{\text{gen}}(C) := \text{Tr}(\rho L_C) + S_{\text{bulk}}(C).$$

Deriving Newton's constant from edge entropy density.

Rather than normalize L_C by fiat, we *derive* the relation to G from the UV edge structure. In the collar double-scaling limit, the edge contribution becomes extensive along the entangling surface $\Sigma = \partial C$:

$$\text{Tr}(\rho L_C) \approx N_\Sigma \cdot \bar{\ell}(t), \qquad \bar{\ell}(t) := \sum_\alpha p_\alpha \log d_\alpha,$$

where N_Σ is the number of UV cut elements covering Σ and $\bar{\ell}(t)$ is the **single-cell edge entropy** from the heat-kernel distribution (Theorem 6.20). Similarly, the geometric area is extensive:

$$A(C) \approx N_\Sigma \cdot a_{\text{cell}},$$

where a_{cell} is the area per UV cut element in the emergent metric.

Matching these expressions gives the **derived formula for Newton's constant**:

$$G = \frac{a_{\text{cell}}}{4\,\bar{\ell}(t)}$$

where:

- a_{cell} is fixed operationally from the UV correlation/mixing length ξ via $a_{\text{cell}} \sim \xi^2$ (from the exponential mixing hypothesis),
- $\bar{\ell}(t) = \sum_R p_R(t) \log d_R$ is computed from the heat-kernel edge distribution with $p_R \propto d_R e^{-t\lambda_R}$.

Explicitly:

$$\bar{\ell}(t) = \frac{\sum_R d_R e^{-t\lambda_R} \log d_R}{\sum_R d_R e^{-t\lambda_R}}.$$

Thus G is the inverse edge-entropy density per geometric area, computable from the UV regulator and the reference-state Gibbs parameter t, rather than a normalization convention.

3.6.5 Entanglement equilibrium from MaxEnt

MaxEnt selection implies that for variations preserving cap labels (fixed size and charges),

$$\delta S_{\text{gen}}(C) = 0.$$

Using the split above and the first law for the bulk term,

$$\delta S_{\text{gen}}(C) = \delta\langle L_C\rangle + \delta\langle K_{\text{bulk}}\rangle.$$

3.6.6 Einstein equation from cap equilibrium

In the EFT regime, combine:

1) Modular energy as stress-tensor charge (Section 5.2 or Section 5.3), and
2) The geometric identity for area variation at fixed volume:

$$\delta A|_{V,\lambda} = -\frac{\Omega_{d-2}\ell^d}{d^2-1}(G_{00} + \lambda g_{00}).$$

The equilibrium condition yields

$$G_{00} + \Lambda g_{00} = 8\pi G\langle T_{00}\rangle,$$

in the diamond rest frame, with Λ fixed by the reference curvature.

3.6.7 Overlaps supply all timelike directions

Different observers through the same bulk point select different local rest frames u. Overlap consistency forces the scalar relation to hold for all timelike u, so

$$G_{ab} + \Lambda g_{ab} = 8\pi G\langle T_{ab}\rangle.$$

3.6.8 Non-tunable numerical constants

The gravity chain yields specific numerical constants as rigid outputs of the axiom chain.

The 2π KMS normalization. From Theorem 4.2, once the OPH refinement branch is geometric, cap modular flow is fixed by the KMS condition to the standard 2π normalization. This is the same rigidity that fixes Unruh/Hawking temperature normalization. The period is determined by the stated scaling-limit package together with the geometric-branch hypothesis.

The geometric coefficient $\Omega_{d-2}/(d^2-1)$. This coefficient appears in both (a) the CFT-ball modular Hamiltonian weight integral and (b) the geometric area-variation identity. It is an exact integral identity:

$$\int_{B_\ell^{d-1}} \frac{\ell^2 - r^2}{2\ell}\, d^{d-1}x = \frac{\Omega_{d-2}\,\ell^d}{d^2-1}.$$

In $d = 4$:

$$\frac{\Omega_2}{4^2-1} = \frac{4\pi}{15} \approx 0.8377580409572781.$$

This is the reason prefactors cancel cleanly when going from $\delta S_{\mathrm{gen}} = 0$ to the Einstein equation (leaving $8\pi G$ with the 2π fixed by Theorem 4.2).

What is predicted. The framework cleanly separates:

- **Non-tunable constants**: 2π (KMS period), $\Omega_{d-2}/(d^2-1)$ (geometric coefficient), the existence of the Einstein form.
- **Micro-dependent constants**: G (Newton's constant) is the conversion between edge entropy and geometric area (a density of edge degrees of freedom per geometric area), which is model-dependent. Λ is fixed by the reference state/constraints. These require specifying the microscopic model.

3.6.9 Quantitative Markov error and controlled corrections

The role of this subsection is to keep three distinct quantities separate:

1. the raw collar conditional mutual information

$$\varepsilon_\delta := I(A_\delta : D_\delta \mid B_\delta)_\omega;$$

2. the constructive recovery error

$$r_{\mathrm{FR}}(\varepsilon_\delta) := 2\sqrt{1 - e^{-\varepsilon_\delta}} \le 2\sqrt{\varepsilon_\delta};$$

3. the exact-Markov replacement error on one fixed faithful collar model,

$$\eta_\delta^{\mathrm{M}} := 4\lambda_*^{-1}\, \delta^{\mathrm{M}}_{A_\delta : B_\delta : D_\delta}(\varepsilon_\delta),$$

   where $\lambda_* > 0$ is the lower spectral bound needed to compare modular Hamiltonians.

The exact identity for the modular defect expectation is unchanged:

$$\langle \Delta K_\delta \rangle_\omega = -I(A_\delta : D_\delta \mid B_\delta)_\omega = -\varepsilon_\delta,$$

with

$$\Delta K_\delta := K_{A_\delta B_\delta D_\delta} - K_{A_\delta B_\delta} - K_{B_\delta D_\delta} + K_{B_\delta}.$$

What changes is the interpretation of later "exact" collar formulas. For each fixed finite-dimensional collar one may choose an exact Markov reference state σ_δ with

$$\|\omega_\delta - \sigma_\delta\|_1 \le \delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta),$$

and a constructive recovered comparison state $\omega_\delta^{\mathrm{rec}}$ with

$$\|\omega_\delta - \omega_\delta^{\mathrm{rec}}\|_1 \le r_{\mathrm{FR}}(\varepsilon_\delta).$$

These two comparisons do different jobs:

> **Bounded-observable control.** For every bounded observable X supported on $A_\delta \cup B_\delta$,
>
> $$|\mathrm{Tr}[X(\omega_\delta - \omega_\delta^{\mathrm{rec}})]| \le \|X\|_\infty\, r_{\mathrm{FR}}(\varepsilon_\delta),$$
>
> and
>
> $$|\mathrm{Tr}[X(\omega_\delta - \sigma_\delta)]| \le \|X\|_\infty\, \delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta).$$
>
> The first bound is constructive and dimension-free; the second is the fixed-collar route to the exact Markov set.
>
> **Modular-additivity control.** If the relevant marginals are uniformly faithful with lower spectral bound $\lambda_* > 0$, then
>
> $$\|\Delta K_\delta(\omega) - \Delta K_\delta(\sigma_\delta)\|_\infty \le \eta_\delta^{\mathrm{M}}.$$
>
> Since $\Delta K_\delta(\sigma_\delta)$ is central by exact Markovity, the finite-stage modular defect is within η_δ^{M} of the exact splice or additivity value.

Therefore the manuscript's exact collar identities arise as limits with a carried remainder

$$\eta_\delta := r_{\mathrm{FR}}(\varepsilon_\delta) + \eta_\delta^{\mathrm{M}},$$

together with the separate long-wavelength derivative remainders present in the small-ball expansion. Exact identities at finite collar width require exact Markovity; otherwise one works with η_δ-controlled approximations and then lets $\eta_\delta \to 0$ in the controlled collar limit.

Modified Einstein relation with an explicit carried remainder. In the small-ball rest-frame calculation, write

$$\delta S_C = \frac{8\pi^2\ell^4}{15}\,\delta\langle T_{00}\rangle + \delta\langle E_C^{(\delta)}\rangle + O(\ell^5\partial T),$$

where $E_C^{(\delta)}$ packages the carried collar remainder. Then the fixed-cap stationarity argument gives

$$\delta(G_{00} + \Lambda g_{00}) = 8\pi G\,\delta\langle T_{00}\rangle + \mathcal{E}_\delta,$$

with

$$\mathcal{E}_\delta := \frac{15G}{\pi\ell^4}\,\delta\langle E_C^{(\delta)}\rangle + O(\ell\,\partial T),$$

and, by the bounds above,

$$|\delta\langle E_C^{(\delta)}\rangle| \le C_C\,\eta_\delta$$

for some collar-dependent constant C_C on the fixed faithful model.

If one wishes to package this carried remainder as an effective anomalous energy density, the natural definition is

$$\delta\langle T_{00}^{\text{anom}}\rangle := \frac{15}{8\pi^2\ell^4}\,\delta\langle E_C^{(\delta)}\rangle.$$

Then

$$\delta(G_{00}+\Lambda g_{00}) = 8\pi G\,\delta(\langle T_{00}\rangle + \langle T_{00}^{\text{anom}}\rangle) + O(\ell\,\partial T).$$

This is a bookkeeping definition of the finite-stage remainder, not an upgrade of the recovered-core theorem. Its significance is exactly that it is controlled:

$$|\delta\langle T_{00}^{\text{anom}}\rangle| \le \frac{15\,C_C}{8\pi^2\ell^4}\,\eta_\delta.$$

The framework carries the Markov error through three distinct controls:

1. $r_{\text{FR}}(\varepsilon_\delta)$ controls bounded observables at one stage;
2. $\delta^{\text{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta)$ controls convergence to the exact Markov normal form on a fixed collar;
3. η_δ^{M} controls modular-Hamiltonian replacements once faithfulness is assumed;
4. the Einstein-branch remainder is a carried $O(\eta_\delta)$ term, not a silent exact identity at finite cutoff.

This is the concrete bridge from "axioms about screens" to "precision GR predictions plus a disciplined carried remainder".

3.6.10 Focusing/QNEC internalization via relative entropy

Once the fixed-cutoff null modular structure of Section 5.2 is combined with the downstream density-upgrade and null-stress premises, focusing constraints follow from information-theoretic principles in the same scaling-limit branch rather than from a separate fixed-cutoff postulate.

Derivation chain. QNEC and focusing are supported conditionally within the same branch:

- Fixed-cutoff null modular bridge (§5.2) $\Longrightarrow$ exact-or-controlled strip additivity and a weak tail generator
- Derived half-sided modular pair + downstream density upgrade $\Longrightarrow$ half-line modular densities and [K, P] = i 2π P
- Relative entropy monotonicity $\Longrightarrow$ QNEC
- Einstein (Thm 5.1) + Raychaudhuri $\Longrightarrow$ QFC for S_gen

Relative entropy monotonicity argument. The key input is the monotonicity of relative entropy under partial trace, which is pure information theory:

$$S(\rho_{AB}\|\sigma_{AB}) \ge S(\rho_A\|\sigma_A).$$

For null deformations parameterized by λ, consider nested null regions $R(\lambda) \subset R(\lambda')$ obtained by varying the entangling cut along v. The modular Hamiltonian K_λ generates the modular flow, and relative entropy satisfies convexity:

$$\frac{d^2}{d\lambda^2}S(\rho_\lambda\|\sigma_\lambda) \ge 0.$$

Proposition 5.10a (Conditional internal QNEC). Under the fixed-cutoff null bridge of §5.2 together with the downstream density-upgrade and resulting null-stress-identification premises, the second null variation of von Neumann entropy satisfies

$$\frac{d^2 S_{\text{bulk}}}{d\lambda^2} \leq 2\pi \langle T_{kk}(\lambda) \rangle,$$

with the 2π normalization fixed by Theorem 4.2.

Proof. Use the derived half-sided modular-inclusion statement of Corollary 5.2e and its Borchers-Wiesbrock consequence from Lemma 5.2f. This gives the translation structure with $[K, P] = -i2\pi P$.

Consider the relative entropy $S(\rho_\lambda \| \omega_\lambda)$ between the state ρ restricted to $R(\lambda)$ and the reference state ω. Monotonicity under restriction to smaller regions ($\lambda' > \lambda$) gives:

$$S(\rho_{R(\lambda)} \| \omega_{R(\lambda)}) \leq S(\rho_{R(\lambda')} \| \omega_{R(\lambda')}).$$

Using the first law $\delta S = \delta \langle K \rangle$ and the geometric half-line modular form supplied by Section 5.2, expand to second order in the deformation. The convexity of relative entropy yields the QNEC inequality. The bound saturates for coherent states in the standard way. QED.

Corollary 5.10b (QFC for generalized entropy). With the central area operator L_C from EC/MaxEnt (Section 5.4), define

$$S_{\text{gen}} = \text{Tr}(\rho L_C) + S_{\text{bulk}}.$$

Given the conditional Einstein branch (Theorem 5.1) and the classical Raychaudhuri identity for null congruences, the Quantum Focusing Conjecture (QFC) follows within the same scaling regime: the generalized expansion Θ_{gen} is non-increasing along null generators.

Proof sketch. The Raychaudhuri equation relates expansion evolution to R_{kk}. Einstein's equation gives $R_{kk} = 8\pi G(T_{kk} - \frac{1}{2} g_{kk} T)$. For null k, this simplifies to $R_{kk} = 8\pi G T_{kk}$. The QNEC (Prop 5.10a) then bounds the bulk entropy production, ensuring $d\Theta_{\text{gen}}/d\lambda \leq 0$. QED.

Significance. This section shows how the Recoverable Generalized Entropy focusing package is internally supported inside the same null-modular branch once the stress-tensor identification and Einstein relation are in place. That package remains part of the stated axiom set.

Theorem 5.1 (Observer-consistency implies a conditional scaling-limit Einstein branch). Under the five OPH axioms, the collar and null-strip hypotheses of Sections 2 and 5, the canonical technical premises **T1**, **T3**, and **T4**, the conditional null-stress identification of Section 5.2, and admissible first variations about a maximally symmetric reference state, the cap equilibrium condition implies the rest-frame first-variation relation

$$\delta(G_{00} + \Lambda g_{00}) = 8\pi G \, \delta \langle T_{00} \rangle.$$

If this relation holds for all local directions and reference states in the locally Lorentzian scaling regime, overlap consistency upgrades it to the semiclassical Einstein equation modulo the expected Λg_{ab} ambiguity.

Proof sketch. The cap first law identifies δS_C with $\delta\langle K_C\rangle$, while the conditional null bridge relates the bulk modular variation to a null stress-tensor charge. Since the reference state is maximally symmetric, the fixed-volume small-ball area identity enters at first order in the perturbation, so entanglement equilibrium yields the displayed rest-frame relation. Overlap consistency then supplies all local directions, and the null-to-tensor lemma upgrades the relation to the tensor equation modulo Λg_{ab}. QED.

3.6.11 Discrete horizon area spectrum and Hawking emission (established)

The edge-sector structure implies a discrete area spectrum with observable consequences for black hole emission. This section develops a established but sharp prediction.

Area eigenvalues from edge sectors. The central area operator (Section 5.4) is

$$L_C = \sum_{\alpha} (\log d_\alpha) P_\alpha,$$

where $d_\alpha \in \mathbb{N}$ is the dimension of the edge Hilbert space in sector α. With the normalization $\mathrm{Tr}(\rho L_C) = \langle A\rangle/4G$, the area eigenvalues are

$$A_\alpha = 4G \log d_\alpha = 4\ell_p^2 \ln d_\alpha,$$

where $\ell_p^2 = \hbar G/c^3$ is the Planck area. Since d_α is a positive integer, **areas are discretely spaced** with logarithmic gaps.

Hawking emission energy quantization. For a Schwarzschild black hole with $A(M) = 16\pi G^2 M^2/c^4$, a transition between sectors $d \to d'$ changes the area by

$$\Delta A = 4\ell_p^2 \ln(d'/d).$$

The corresponding energy at infinity is $\Delta E = c^2 \Delta M$, with $\Delta M = \Delta A/(dA/dM)$. This gives

$$\Delta E = \frac{\hbar c^3}{8\pi G M} \ln(d'/d).$$

Using the Hawking temperature $T_H = \hbar c^3/(8\pi G k_B M)$, whose 2π normalization is fixed by the BW_{S^2} tangent-limit normalization of Theorem 4.2:

$$\Delta E = k_B T_H \ln(d'/d).$$

Integer transitions. If dominant transitions multiply the edge dimension by an integer k (i.e., $d'/d = k$), the spectrum becomes a discrete comb:

$$\Delta E_k = k_B T_H \ln k, \qquad \Delta f_k = \frac{c^3}{16\pi^2 G M} \ln k.$$

Structural condition: comb vs. generic discreteness. The log-integer *comb* structure requires the additional dynamical assumption that integer-multiplication transitions (d → kd) dominate. If generic transitions between arbitrary integers dominate instead, the set of ln(d'/d) values becomes a dense log-rational set that may appear quasi-continuous after folding in linewidths and astrophysical effects. What follows directly from the axioms is the *discrete area spectrum*; the clean comb pattern is conditional on the selection rule.

Continuation-level template (Discrete Hawking spectrum). If the additional integer-transition selection rule is realized, the Hawking emission spectrum consists of discrete lines with spacing $\Delta E_k = k_B T_H \ln k$, where k is an integer characterizing the dominant sector transitions, instead of a continuous thermal profile.

Mass-independent fractional linewidth. Using Page's semiclassical calculation for emission power $P(M) = p_0 \hbar c^6/(G^2 M^2)$ with $p_0 \approx 2 \times 10^{-4}$, the emission rate is $\dot{N} \approx P/\langle E \rangle$ where $\langle E \rangle = a\, k_B T_H$ with $a \sim \mathcal{O}(1-10)$. The natural linewidth $\Gamma \sim \hbar \dot{N}$ divided by the level spacing gives:

$$\frac{\Gamma}{\Delta E_k} \approx \frac{64\pi^2 p_0}{a \ln k} \approx 3 - 5\%$$

independent of black hole mass. Within this continuation branch, the emission lines are narrow (few-percent fractional width) and the fraction is mass-independent.

Connection to quasinormal modes. The highly-damped Schwarzschild quasinormal modes have asymptotic real part (Motl, 2002):

$$\mathrm{Re}\,\omega \to \frac{c^3}{8\pi G M} \ln 3.$$

This matches exactly the $k = 3$ transition frequency $\Delta E_3/\hbar$. If one adopts a Bohr-type identification between quantum transition frequencies and asymptotic QNM frequencies, this selects

$$\Delta A = 4\ell_p^2 \ln 3 \approx 4.39\, \ell_p^2$$

as the fundamental area quantum.

Conditionality statement. The area quantization follows from the edge-sector structure (derived). The $k = 3$ selection requires the additional interpretive identification with QNM frequencies (not derived from axioms). The linewidth prediction uses standard semiclassical inputs.

Numerical examples. For Δf_k = ($c^3/16\pi^2$GM) ln k:

- **M = 30 M⊙**: k=2 at 29.7 Hz, k=3 at 47.1 Hz
- **M = 1 M⊙**: k=2 at 891 Hz, k=3 at 1412 Hz
- **M = 10^{12} kg (primordial)**: k=2 at 7.3 MeV, k=3 at 11.6 MeV

These frequencies track $k_B T_H \ln k$ exactly and are in principle distinguishable from a continuous thermal spectrum.

Experimental test: PBH burst searches with comb template.

The discrete Hawking comb provides an OPH-unique signature that can be tested against existing gamma-ray data. The smoking gun is **log-integer energy ratios**: if two emission lines are observed at energies E_2 and E_3, their ratio must satisfy

$$\frac{E_3}{E_2} = \frac{\ln 3}{\ln 2} \approx 1.585$$

exactly, independent of black hole mass. This is a parameter-free prediction.

Available instruments and energy coverage. The $k = 2$ line energy $E_2 = k_B T_H \ln 2$ determines which instruments can see a given BH mass:

Instrument	Energy band	BH mass range (k=2 in band)
Fermi GBM (BGO)	0.15–40 MeV	$2{\times}10^{11}$–$5{\times}10^{13}$ kg
Fermi LAT	0.1–300 GeV	$2{\times}10^{7}$–$7{\times}10^{10}$ kg
H.E.S.S.	0.1–100 TeV	$7{\times}10^{4}$–$7{\times}10^{7}$ kg
LHAASO-WCDA	1–15 TeV	$5{\times}10^{5}$–$7{\times}10^{6}$ kg

Instrument	Energy band	BH mass range (k=2 in band)

Detector resolution vs. intrinsic linewidth. The predicted intrinsic linewidth is 3–5% (mass-independent). Current detector energy resolutions:

- Fermi GBM: $< 10\%$ (0.1–1 MeV), $\sim 4\%$ at 10 MeV (BGO)
- Fermi LAT: $< 10\%$ (1–100 GeV)
- H.E.S.S.: $\sim 15\%$ (TeV)
- LHAASO-WCDA: $\sim 33\%$ (TeV)

The comb is in principle resolvable with GBM/LAT; at TeV energies it would appear as moderately broad bumps rather than sharp lines.

Search protocol. A dedicated OPH-comb search would:

1. Select burst-like candidates (10–120 s time windows, matching existing PBH burst search protocols).
2. Fit each candidate with null model (smooth continuum) vs. OPH comb model (peaks at $E_k = E_0 \ln k$ convolved with detector response).
3. Scan over the single scale parameter $E_0 = k_B T_H$ (equivalently, BH mass).
4. Require at least two lines satisfying log-integer ratio to claim detection.
5. Correct significance for trials (time windows × sky positions × E_0 scan).

Observational context. Dedicated PBH burst searches (H.E.S.S., LHAASO) report **no significant bursts**. An OPH-specific comb-template analysis of archival data would:

- Set upper limits on OPH-comb PBH burst rates
- Demonstrate direct testability of the discrete spectrum prediction
- Provide constraints comparable to or stronger than generic PBH burst limits

Data availability. Fermi GBM provides public Time-Tagged Event (TTE) burst data; Fermi LAT provides public photon event lists with documented analysis workflows. H.E.S.S. has a small public test data release.

GW horizon spectroscopy: conditional continuation template for Kerr remnants.

The same continuation-level discrete-horizon branch extends to gravitational wave observables. For Kerr black holes, the thermodynamic first law is $\delta M = T_H \delta S + \Omega_H \delta J$, so the entropy change for absorbing a quantum with frequency ω and azimuthal number m is:

$$\delta S = \frac{\hbar(\omega - m\Omega_H)}{k_B T_H}.$$

In the edge-sector framework, $\delta S = \ln(d'/d)$, so the discreteness condition becomes:

$$\hbar(\omega - m\Omega_H) = k_B T_H \ln k, \qquad k \in \{2, 3, 4, \ldots\}$$

Under those additional discrete-horizon premises, this gives the **GW horizon spectroscopy comb**: discrete resonant frequencies where the horizon can efficiently absorb or emit energy.

Kerr line frequencies. For a remnant with mass M and dimensionless spin $\chi = a_*/M$, define the spin correction factor:

$$g(\chi) = \frac{2\sqrt{1-\chi^2}}{1+\sqrt{1-\chi^2}}, \qquad \Omega_H(M,\chi) = \frac{c^3}{2GM} \cdot \frac{\chi}{1+\sqrt{1-\chi^2}}.$$

The line frequencies are:

$$f_{k,m}(M,\chi) = \frac{m\,\Omega_H(M,\chi)}{2\pi} + \frac{c^3}{16\pi^2 GM}\, g(\chi)\, \ln k$$

This is rigidly constrained: once LIGO/Virgo infers (M, χ) for a remnant, the entire line pattern is fixed with no free parameters.

Line weights from GR envelope + discretization. The line strengths are not arbitrary; they are fixed by matching to the known GR greybody absorption spectrum in the semiclassical limit. The discretization rule gives bin width $\Delta\omega_k \approx \omega_T \ln(1 + 1/k)$ where $\omega_T = k_B T_H/\hbar$. The net line weight (absorption minus stimulated emission) is:

$$W^{\rm net}_{k,\ell m} = \Gamma^{\rm GR}_{\ell m}(\omega_{k,m}) \cdot \Delta\omega_k \cdot \frac{k-1}{k}$$

where $\Gamma^{\rm GR}_{\ell m}$ is the standard GR greybody factor and the $(k-1)/k$ factor arises from KMS detailed balance with $e^{(\omega - m\Omega_H)/T_H} = k$.

Universal stacking coordinate. Define the dimensionless rescaled frequency:

$$x := \frac{GM}{c^3 g(\chi)}(\omega - m\Omega_H).$$

Then the predicted line locations collapse to universal constants:

$$x_k = \frac{\ln k}{8\pi} \qquad (k = 2, 3, 4, \ldots)$$

Numerically: $x_2 = 0.02758$, $x_3 = 0.04371$, $x_4 = 0.05516$, $x_5 = 0.06404$.

Stacking test. Multiple BBH events can be mapped to this universal x coordinate and stacked. If the comb is real, peaks align across events with different (M, χ); detector noise does not stack coherently.

Comparison to existing work. Prior area-quantization searches [36] used parameterized models with one free spacing constant. The OPH prediction is more constrained: multiple lines with exact $\ln k$ ratios, plus the $(k-1)/k$ weight hierarchy from detailed balance.

Numerical example (GW170608). Remnant parameters: $M_f \approx 18.0 M_\odot$, $\chi_f \approx 0.69$. For $m = 2$, the horizon rotation frequency is $m\Omega_H/(2\pi) \approx 719$ Hz. The **thermal comb spacing** (the part that encodes the area quantization) is:

k	$\Delta f_k := \frac{c^3 g(\chi)}{16\pi^2 GM} \ln k$ (Hz)	Relative weight $(k-1)/k$
2	41.6	0.500
3	65.9	0.667
4	83.2	0.750
5	96.5	0.800
6	107.5	0.833

The full physical frequencies are $f_{k,2} = 719 + \Delta f_k$ Hz (i.e., 760–827 Hz), outside LIGO's most sensitive band for this remnant. However, the **stacking analysis** uses the rescaled coordinate

$x = GM(\omega - m\Omega_H)/(c^3 g(\chi))$, which maps the thermal spacing to universal constants $x_k = \ln k/8\pi$ regardless of the rotation offset.

Measurement contradiction criterion. The smoking gun is the rigid arithmetic pattern: after rescaling by (M, χ), spectral features must satisfy $f_k/f_2 = \ln k/\ln 2$ exactly, independent of remnant parameters. Absence of coherent stacking at the predicted x_k values identifies a measurement contradiction with the log-integer area spectrum.

3.6.12 Classical mechanics from emergent GR

Once the Einstein equation is established, the framework inherits standard GR consequences. This section makes explicit how classical mechanics emerges.

Stress-energy conservation is automatic. The contracted Bianchi identity is geometric:

$$\nabla^a G_{ab} = 0.$$

Combined with the Einstein equation, this implies:

$$\nabla^a \langle T_{ab} \rangle = 0.$$

Geodesic motion from dust limit. For pressureless classical matter ("dust"), $T^{ab} = \rho\, u^a u^b$. Conservation yields:

$$\nabla_a(\rho u^a u^b) = 0 \quad \Rightarrow \quad u^b \nabla_a(\rho u^a) + \rho\, u^a \nabla_a u^b = 0.$$

Projecting orthogonally to u^b using $h^b{}_c = \delta^b{}_c + u^b u_c$ kills the first term, giving:

$$\rho\, u^a \nabla_a u^b = 0 \quad \Rightarrow \quad u^a \nabla_a u^b = 0.$$

This is the geodesic equation: free classical bodies follow spacetime geodesics.

Newtonian limit from weak-field GR. Take the weak-field, slow-motion limit with metric:

$$g_{00} \approx -(1 + 2\Phi/c^2), \qquad g_{0i} \approx 0, \qquad g_{ij} \approx \delta_{ij}(1 - 2\Phi/c^2),$$

and velocities $|\mathbf{v}| \ll c$. Then $G_{00} \approx 2\nabla^2\Phi/c^2$ (leading order), and $T_{00} \approx \rho c^2$. The Einstein equation reduces to:

$$\nabla^2 \Phi = 4\pi G \rho.$$

Geodesic motion reduces to:

$$\ddot{\mathbf{x}} = -\nabla\Phi.$$

These are Newton's gravitational law and Newton's second law. Classical mechanics is recovered as a controlled limit of the emergent GR dynamics.

Precision classical predictions. Once the field equation is fixed to Einstein form, the framework inherits the standard GR precision toolbox (post-Newtonian expansion, lensing, time delay, etc.), with no free "shape" parameters beyond G and Λ.

Selected precision predictions (in the regime where the GR derivation applies):

Light bending by mass M: For impact parameter b,

$$\Delta\theta = \frac{4GM}{c^2 b}.$$

For the Sun with $b \approx R_\odot$: $\Delta\theta \approx 1.751$ arcsec.

Mercury perihelion advance: Per orbit,

$$\Delta\varpi = \frac{6\pi GM}{a(1-e^2)c^2}.$$

Using Mercury's orbital parameters: $\Delta\varpi \approx 42.98$ arcsec/century.

Gravitational redshift: Between two radii in a static potential,

$$\frac{\Delta\nu}{\nu} \approx \frac{\Delta\Phi}{c^2}.$$

For the Sun (surface to infinity): $z \approx 2.12 \times 10^{-6}$.

These predictions are fixed functions of G and known source parameters, and are confirmed observationally to high precision. The framework inherits them automatically once the Einstein equation is derived.

3.6.13 Precision gravity predictions and experimental bounds

The gravity sector makes symmetry-protected exact-zero predictions that can be confronted with the tightest available experimental bounds. This section translates the theoretical predictions into the specific observables that experiments actually constrain.

Speed of gravitational waves. The derived GR regime implies massless gravitons propagating on the same null cones as photons:

$$\frac{c_{\mathrm{GW}} - c}{c} = 0 \text{ exactly.}$$

Current bound (GW170817 + GRB 170817A multi-messenger):

$$-3 \times 10^{-15} < \frac{c_{\mathrm{GW}} - c}{c} < +7 \times 10^{-16} \quad \text{(90\% credibility).}$$

For a source at ~ 40 Mpc, this fractional difference corresponds to only a few seconds of propagation-time mismatch across $\sim 10^8$ years of travel.

Graviton mass. The gauge redundancy (diffeomorphism invariance) forbids a hard mass term:

$$m_g = 0 \text{ exactly.}$$

Current bound (GW dispersion analysis, PDG 2025):

$$m_g \leq 1.76 \times 10^{-23} \text{ eV}/c^2 \quad \text{(90\% credibility).}$$

This corresponds to a reduced Compton wavelength $\bar{\lambda}_C \gtrsim 1.6 \times 10^{16}$ m, i.e., order ~ 1.6 light-years.

No dipole radiation. Many modified gravity theories predict extra channels (scalar/vector) producing dipolar radiation at (-1)PN order. The derived GR limit predicts no such channel.

Current bound (GW170817 inspiral phasing, PDG 2025):

$$-4 \times 10^{-6} < \delta\hat{p}_{-2} < 2 \times 10^{-5} \quad \text{(90\% credibility).}$$

Only tensor polarizations. The GR outcome means only the two tensor (helicity-2) modes propagate. Pure non-tensor hypotheses are disfavored by observational constraints, and mixed tensor-scalar/vector models are tightly constrained.

Equivalence principle tests. Additional null checks from the derived GR structure:

- Universality of free fall (space tests): precision ~10^{-15}
- Nordtvedt parameter ($\eta = 4\beta$ - γ): (0.47 ± 0.55) × 10^{-4}
- Binary pulsar radiative damping (PSR J0737-3039): 0.999963 ± 0.000063

3.6.14 Theory-side error propagation from Markov bounds

The framework provides exact-zero predictions and quantitative control over how well those predictions hold. The Markov/recovery machinery can be propagated through the entire GR emergence chain.

The key quantitative hook. From Theorem 3.1, if the target state satisfies

$$I(A_k : C_k \mid B_k) \le \varepsilon_k,$$

then recovery maps exist with trace-distance error

$$\delta_k = 2\sqrt{\ln 2 \cdot \varepsilon_k}.$$

Trace distance gives immediate bounds on observable errors. Using the standard dual norm inequality:

$$|\langle O\rangle_\rho - \langle O\rangle_\sigma| \le \|O\|_\infty \|\rho - \sigma\|_1 = 2\|O\|_\infty D(\rho, \sigma),$$

where $D(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$ is the trace distance.

Exponential decay from the mixing hypothesis. The mixing assumption (Section 2.3) provides:

$$I_\omega(A_\delta : D_\delta \mid B_\delta) \le c \cdot |\partial C|_{\mathrm{UV}} \cdot e^{-\delta/\xi}.$$

Combining these gives an explicit precision dial:

$$\delta_{\mathrm{step}} \lesssim 2\sqrt{\ln 2 \cdot c \cdot |\partial C|_{\mathrm{UV}}} \cdot e^{-\delta/(2\xi)}.$$

What precision requires. To match the GW speed bound ($\sim 10^{-15}$ fractional accuracy), the recovery-map error must satisfy:

$$\delta \lesssim 10^{-15} \quad \Rightarrow \quad \varepsilon \lesssim \frac{(\delta/2)^2}{\ln 2} \approx 3.6 \times 10^{-31}.$$

This is extremely small, but achievable: with a macroscopic boundary ($|\partial C|_{\mathrm{UV}} \sim 10^{35}$ for a meter-scale boundary at Planck UV scale), the exponential decay $e^{-\delta/\xi}$ with $\delta/\xi \sim$ a few hundred easily pushes below 10^{-31} once the prefactor is included.

Precision summary. The framework provides:

1. Exact-zero predictions ($m_g = 0$, $c_{\mathrm{GW}} = c$) from symmetry protection.
2. Translation of those zeros into the specific observables experiments constrain.
3. Explicit bounds on how far derived geometric statements can drift, using the conditional mutual information → trace distance → observable error chain.

This is the concrete path from "axioms about screens" to "precision GR predictions with quantitative error control."

3.6.15 Dark-sector response from the modular anomaly

This subsection studies the benchmark scale obtained by adding a specific deep-IR response ansatz on top of the modular-anomaly term. Neither the observational dark-matter identification nor MOND/RAR-like dynamics are derived here.

The modular anomaly term T_{ab}^{anom} derived in Section 5.9 supplies one structural ingredient for a possible dark-sector continuation, without introducing new particle species.

The identification. The anomalous stress-energy contribution

$$\langle T_{00}^{\text{anom}} \rangle = \frac{15}{8\pi^2} \cdot \frac{\delta \langle K_C^{(\text{anom})} \rangle}{\ell^4}$$

is "dark" by construction: it arises from information-theoretic/gravitational structure (modular Markov imperfections), not from Standard Model fields. It gravitates but does not couple electromagnetically. These properties explain why the term is explored as a dark-sector ingredient, but they do not by themselves close the observational identification.

Connection to the cosmological constant. The framework makes Λ a global capacity parameter:

$$\Lambda = \frac{3\pi}{G N_{\text{scr}}}, \quad r_{dS} = \sqrt{\frac{3}{\Lambda}}.$$

This introduces an unavoidable IR length scale r_{dS}. Any galaxy-scale continuation built from the anomaly term would therefore be an IR phenomenon, appearing only when accelerations are small and distances are large.

Acceleration benchmark. If one assumes the relevant deep-IR response is controlled only by r_{dS}, c, and the fixed anomaly prefactor, then the benchmark scale must:

1. Vanish if $r_{dS} \to \infty$ (infinite capacity, no de Sitter scale)
2. Be controlled by r_{dS} as the only new IR scale
3. Carry non-tunable coefficients from the derivation

The anomaly enters with prefactor $\frac{15}{8\pi^2}$. The corresponding benchmark acceleration scale constructible from (Λ, c) is:

$$\boxed{a_0^{(\text{OPH})} := \frac{15}{8\pi^2} \cdot c^2 \sqrt{\frac{\Lambda}{3}} = \frac{15}{8\pi^2} \cdot \frac{c^2}{r_{dS}}}$$

Normalization estimate. Using Planck 2018 ΛCDM parameters ($H_0 \approx 67.4$ km/s/Mpc, $\Omega_\Lambda \approx 0.685$):

- $\Lambda \approx 1.09 \times 10^{-52}\ \text{m}^{-2}$
- $r_{dS} \approx 1.66 \times 10^{26}$ m
- Therefore:

$$\boxed{a_0^{(\text{OPH})} \approx 1.03 \times 10^{-10}\ \text{m/s}^2}$$

For comparison, observational fits to galaxy regularities (RAR/MDAR/MOND phenomenology) quote $a_0 \sim 1.2 \times 10^{-10}\ \text{m/s}^2$. This numerical proximity does not by itself derive galaxy dynamics.

One illustrative response ansatz. If one further assumes that T_{00}^{anom} is the dominant deep-IR source and that the response organizes into a MOND/RAR-like law, the Newtonian limit could be written as:

$$\nabla^2\Phi = 4\pi G(\rho_b + \rho_{\mathrm{anom}}),$$

i.e., baryons plus an effective extra density. Under the same extra ansatz the radial acceleration relation (RAR) could be written as:

$$g_{\mathrm{obs}} \approx g_b + \sqrt{a_0 \cdot g_b}, \quad g_{\mathrm{DM}} := g_{\mathrm{obs}} - g_b \approx \sqrt{a_0 \cdot g_b}.$$

With $a_0 = a_0^{(\mathrm{OPH})}$ fixed, that same ansatz would imply:

(i) Baryonic Tully-Fisher relation.

$$V^4 \approx G \cdot M_b \cdot a_0^{(\mathrm{OPH})}$$

where V is the asymptotic rotation velocity and M_b is baryonic mass.

(ii) Flat rotation curves. For a point mass M_b:

$$g_{\mathrm{DM}}(r) = \frac{\sqrt{GM_b\, a_0^{(\mathrm{OPH})}}}{r} \quad \Rightarrow \quad M_{\mathrm{DM}}(r) \propto r$$

i.e., inferred dark mass grows linearly with radius, producing flat rotation curves.

(iii) Characteristic surface density.

$$\Sigma_0^{(\mathrm{OPH})} = \frac{a_0^{(\mathrm{OPH})}}{2\pi G} \approx 0.25\ \mathrm{kg/m^2} \approx 120\, M_\odot/\mathrm{pc}^2.$$

This lies in the range of observed central halo surface densities, but it is not a derived halo theorem.

Status. What is grounded in the framework developed here:

- The modular anomaly term exists with fixed coefficient $\frac{15}{8\pi^2}$
- Λ and r_{dS} are determined by screen capacity

What remains missing for theorem-level closure:

- A controlled nonrelativistic limit from the anomaly term to galaxy observables
- A derived response law selecting the MOND/RAR functional rather than alternative IR behavior
- Lensing, cluster, and Bullet-Cluster phenomenology
- Cosmological abundance and structure-formation checks
- Environment-dependence and stability control

This branch contains structural ingredients and an IR benchmark, but it does not yet derive the precise galaxy-scale response from the recovered core.

Falsifiability. If later work cannot supply the missing derivation steps, or if any resulting continuation is incompatible with galaxy, lensing, cluster, Bullet-Cluster, or cosmological data, then the modular-anomaly dark-sector continuation retracts. The recovered OPH core does not.

3.6.16 De Sitter holography: static patch vs boundary-at-infinity

A natural question arises: how does this framework relate to the "unsolved problem" of de Sitter holography?

What the usual dS holography problem is. When people say "dS holography is unsolved," they typically mean that we do not have anything as sharp as AdS/CFT, where the bulk has a timelike asymptotic boundary supporting a well-defined dual CFT with a precise dictionary. For de Sitter, there is no asymptotic timelike boundary in the static patch where one can simply place the dual theory. The classic dS/CFT proposal at future infinity has familiar difficulties, including non-unitarity worries and complex conformal weights.

What this model does differently. The framework begins with an observer's static patch and its horizon screen S^2, building a net of subregion algebras on that screen. At finite cutoff those algebras are type-I regulators. The Lorentz branch, when invoked, is a scaling-limit statement about the refinement-limit observer net, and that limit may leave the regulator class. On the **T1** branch of Theorem 3.5.2, the realized scaling-limit cap modular action is geometric and, in the non-type-I case of interest, generally outer. That branch identification is *not* derived from MaxEnt alone; it is exactly the explicit **T1+T3** boundary.

This is therefore a fundamental fork away from AdS/CFT-style holography:

AdS/CFT	This framework
Codimension-1 boundary at infinity	Codimension-2 horizon screen (S^2)
Single global boundary theory	Observer-dependent patches that overlap
Dual CFT required	Only algebras + consistency conditions
Negative Λ	Positive Λ natural

This aligns with the static-patch/complementarity intuition in the dS literature, where the fundamental description is patch-based and different static patches are related by consistency rules, not by a single global boundary theory.

The mechanism: Λ as global capacity, not local physics. A key structural result is that null modular data reconstruct the stress tensor only up to an additive metric term. This is the statement that vacuum-energy or cosmological-constant shifts are invisible to the local null-data route. The Einstein equation derived from entanglement equilibrium is therefore fixed only up to Λg_{ab}.

For that reason Λ cannot be determined by local consistency alone. It must be fixed by a *global* input: the total number of degrees of freedom on the screen. The de Sitter link is

$$\Lambda = \frac{3\pi}{G N_{\mathrm{scr}}}.$$

If the screen Hilbert space has finite dimension, the natural semiclassical interpretation is a cosmological horizon with finite entropy. The local/global stack is therefore: local null-modular reconstruction leaves the metric ambiguity, global screen capacity fixes the remaining constant. This logic is compatible with the BW scaling branch but does not depend on the scaling-limit algebra remaining type I.

What this solves vs. what it assumes. The model does **not** solve the classic "give me a unitary CFT at future infinity" problem. It does not aim there. It also does **not** prove that every refinement-stable MaxEnt branch lands in the BW/canonical static-patch phase. What it does provide is a coherent route to **patch holography** in which de Sitter static patches are natural:

1. the fundamental object is a horizon screen in a static-patch description;
2. Λ is a capacity parameter tied to finite Hilbert-space dimension, not a locally reconstructible vacuum-energy term;
3. Einstein-like dynamics emerge up to Λg_{ab};
4. on the explicit **T1+T3** branch, the realized scaling-limit cap modular action is geometric and may be outer on a non-type-I observer algebra.

BW-side status. The Lorentz side stays at the explicit conditionality boundary recorded elsewhere in the manuscript. The internal extension route is a two-object scaffold: first a realized scaling-limit cap-pair extraction from transported marginals, then ordered cut-pair rigidity on that realized limit. The pressure point is the first object. Until that realized scaling-limit cap pair exists, the BW/canonical language remains theorem-external shorthand for the **T1+T3** branch and does not become a fully internalized static-patch theorem.

Many observers, one Λ. The philosophical stance of the framework maps naturally onto static-patch intuition: each timelike observer has a horizon and a patch, and there is no operational access to a single global description. The de Sitter parameter Λ is the one thing that all overlap-consistent descriptions share as a global capacity constraint.

Summary. The model gets de Sitter by moving the holographic screen from "infinity" to an observer's horizon and by elevating de Sitter entropy, i.e. finite screen capacity, to a fundamental input. The usual dS-holography obstacles are precisely the ones avoided by refusing the boundary-at-infinity viewpoint. This is not a claim of a solved dS/CFT dual. It is a static-patch holography program with an explicit local/global split and an explicit BW-branch conditionality boundary.

3.7 Gauge Reconstruction and Standard Model Structure

3.7.1 Edge sector category and gauge group reconstruction

At any fixed UV cutoff, edge-center completion provides finite tensor categories of transportable edge charges. The local MaxEnt / collar-mixing package established earlier controls only fixed-cutoff recoverability, modular-support localization, and carried error terms on the realized branch. It is logically separate from the question whether transportable edge sectors survive refinement, and it supplies no theorem that the refinement-limit sector category is trivial or nontrivial. Whenever compact gauge reconstruction is invoked, that transportable refinement-directed sector system is exactly the content of **T7**.

Assume therefore that in the EFT regime there is a directed system $(\mathsf{Sect}_r)_r$ of transportable edge-sector categories with refinement functors, and write

$$\mathsf{Sect}_\infty := \varinjlim_{r} \mathsf{Sect}_r,$$

for the directed colimit retaining the sectors and intertwiners that persist in that transportable system, with objectwise finite-dimensional fibers. The theorem below is neutral about whether Sect_∞ is trivial or nontrivial: if **T7** supplied only the tensor unit, the reconstructed compact group would be the trivial group.

This gives a tensor category Sect_∞ of edge charges:

- objects: transportable sector labels that persist in the assumed directed system,

- morphisms: intertwiners between sectors,
- tensor product: fusion by collar concatenation, $\alpha \otimes \beta = \bigoplus_\gamma N_{\alpha\beta}{}^{\gamma}\,\gamma$,
- duals: orientation reversal $\alpha \leftrightarrow \bar{\alpha}$,
- symmetric braiding in the EFT regime (no anyonic statistics in 3+1D).

Let $\mathcal{F} : \mathsf{Sect}_\infty \to \mathsf{Hilb}_{\mathrm{fd}}$ be the fiber functor that sends each sector to its edge multiplicity space. The fibers are finite-dimensional objectwise even though Sect_∞ can have infinitely many simple objects. The theorem begins only once that transportable colimit and its fiber functor are given; those data are not consequences of Axiom 3 alone.

Theorem 6.1 (Tannaka/DR reconstruction). If Sect_∞ is a rigid symmetric C^* tensor category with a faithful fiber functor $\mathcal{F}$, then there exists a compact group G, unique up to isomorphism, such that $\mathsf{Sect}_\infty \simeq \mathrm{Rep}(G)$. Moreover,

$$G = \mathrm{Aut}_\otimes(\mathcal{F})$$

is a compact subgroup of a product of unitary groups.

Proof sketch. Define G as the group of monoidal natural automorphisms of $\mathcal{F}$. This is compact because it is closed in a product of unitary groups. By Tannaka-Krein/DR reconstruction, objects and morphisms of Sect_∞ are recovered as finite-dimensional representations and intertwiners of G. The MaxEnt/local-Gibbs/collar-mixing package is not used in this step to prove existence or nontriviality of the transportable colimit; it only supplies the earlier fixed-cutoff collar control. QED.

Corollary 6.1 (field algebra reconstruction, conditional). If in the small-region limit the edge sectors are localized and transportable in the DHR sense (i.e., charges can be moved between patches without changing their fusion), then there exists a field algebra $\mathcal{F}$ and a compact group G such that $\mathcal{A} = \mathcal{F}^G$. This is the Doplicher–Roberts reconstruction of local gauge symmetry from sectors. QED.

Proposition 6.1a (Transportability from gluing obstruction on the central branch). On the central-defect branch, DHR transportability is not an independent assumption. In the gluing framework (Section 3), transportability is precisely the statement that charges can be moved between patches without changing fusion rules. The obstruction to path-independent transport is the Čech-type central cocycle z_{ijk} from the central-defect subbranch.

Explicitly: the gluing framework gives a gauge-invariant Čech 2-cocycle class $[z]$ built from triple overlaps (Section 6.6). Transportability holds iff this class vanishes:

$$\text{DHR transportable} \iff [z] = 0 \iff \text{loop-coherent gluing.}$$

Proof. Transportability means charges can be moved along any path without affecting the result. In gluing language, this is path-independent parallel transport of edge labels. Lemma 6.12 shows that loop-coherent global gluing exists iff $[z] = 0$. But loop-coherent gluing is exactly path-independent transport, so the equivalence holds. QED.

Corollary. On the central branch, the "DHR transportability" condition in Corollary 6.1 is internal to the gluing framework: it is equivalent to requiring that the central obstruction class vanishes. This is a constraint on the allowed sector structure, not an external physical assumption. On the genuinely noncentral branch, Corollary 3.4a gives the parallel criterion in terms of the higher-gauge class q_Σ.

3.7.2 Selecting the SM factors (derived from MAR)

Theorem 6.1 yields *some* compact G. Axiom 5 (MAR, Minimal Admissible Realization) uniquely determines which G is realized.

Axiom 5 (MAR): Minimal Admissible Realization. Among all OPH-realizable sector packages $\mathfrak{S}$ consisting of the connected Lie gauge-sector image relevant in the low-energy EFT, its admissible light chiral matter content, and one Higgs doublet, and which are (i) loop-coherent / transportable (vanishing relevant obstruction: $[z] = 0$ on the central branch or $q_\Sigma = 0$ on the higher-gauge branch), (ii) anomaly-free, (iii) refinement-stable with light chiral matter, (iv) single-Higgs Yukawa-completable with one connected abelian charge factor acting nontrivially on the coupled carrier, (v) intrinsically CP-capable, (vi) weak-sector UV-completable, the realized low-energy package is the lexicographically minimal one under

$$C(\mathfrak{S}) = (\chi_{\mathrm{cpl}},\ N_{\mathrm{nonab}},\ N_c,\ N_g).$$

The complete formal statement, admissibility definitions, and proof route are given in the rest of Section 6.2 and in the dedicated gauge-derivation surface.

Here χ_{cpl} is the **coupled edge capacity**: the dimension of the minimal unitary carrier containing a common irreducible block on which the admissible pseudoreal and complex nonabelian charge types both act nontrivially. This is intentionally stronger than the abstract minimal faithful representation dimension. For $S(U(3) \times U(2))$, the block-diagonal action on $\mathbb{C}^3 \oplus \mathbb{C}^2$ is faithful of dimension 5, but it is not coupled and therefore does not enter MAR. The object minimized by MAR is the sector package $\mathfrak{S}$, not the bare tensor category by itself. MAR is thus an explicit structural-economy selector on the admissible class, not a theorem derived from the earlier axioms.

Note on transport obstruction. The loop-coherent gluing condition is an explicit premise of the gauge derivation. On the central branch this is $[z] = 0$ (Proposition 6.1a); on the genuinely noncentral branch it is $q_\Sigma = 0$ (Corollary 3.4a). In either form it ensures that the reconstructed compact group acts as a genuine global gauge symmetry.

What MAR derives. Product structure, minimal sector content, and coupled edge-capacity minimality are consequences of MAR applied to the admissible class:

- **Product structure**: follows from the minimal coupled carrier $\mathbb{C}^3 \otimes \mathbb{C}^2$, which enforces commuting color and weak actions.
- **Minimal sector content**: the pseudoreal doublet and complex triplet are the minimal nonabelian representations satisfying the admissibility conditions, while the connected abelian charge factor is an explicit admissibility input.
- **Coupled edge-capacity minimality** (formerly S): is the first component of MAR's complexity vector.

With MAR stated, the SM derivation proceeds via standard lemmas:

Lemma 6.2 (Product factorization implies product group). If $\mathsf{Sect} \simeq \mathrm{Rep}(G)$ and $\mathsf{Sect} \simeq \mathsf{Sect}_1 \boxtimes \mathsf{Sect}_2$, then

$$G \cong G_1 \times G_2, \qquad \mathsf{Sect}_i \simeq \mathrm{Rep}(G_i).$$

QED.

Lemma 6.3 (SU(2) from a pseudoreal doublet). Let H be a positive-dimensional compact connected Lie group with a faithful 2D pseudoreal unitary representation V. Then the semisimple part of H contains an $SU(2)$ factor acting as the fundamental doublet. Finite or disconnected counterexamples are not part of the theorem package, which only applies the lemma to the identity component on the relevant nonabelian image. QED.

Lemma 6.4 (SU(3) from an irreducible triplet). Let H be a positive-dimensional compact connected Lie group with a faithful irreducible complex 3D unitary representation W. Then the semisimple image contains an $SU(3)$ factor acting as the fundamental triplet. Finite or disconnected counterexamples are not part of the theorem package, which only applies the lemma to the identity component on the relevant nonabelian image. QED.

Lemma 6.5 (Connected abelian factor criterion). If the admissible sector package contains a connected abelian charge factor acting nontrivially on the coupled carrier, then the identity component of the abelianized reconstructed group contains a one-torus. Under the single connected abelian-factor admissibility premise, this factor is $U(1)$. QED.

Proposition 6.6 (physical group quotient). If the realized matter spectrum has hypercharges quantized in sixths, then the kernel acting trivially on all realized sectors is Z_6, so

$$G_{\text{phys}} = \frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}.$$

QED.

Proposition 6.6a (SM from MAR). Under Axiom 5 (MAR):

- Admissibility conditions (iii)–(iv) require both a pseudoreal nonabelian charge type and a complex nonabelian charge type, and include one connected abelian charge factor acting nontrivially on the coupled carrier (Lemma 6.7, Corollary 6.8).
- The minimal faithful pseudoreal representation is the doublet ($\chi = 2$), giving $SU(2)$. No intrinsically complex 2D irrep qualifies in the connected compact Lie class, because the connected derived image of any irreducible 2D unitary representation lies in $SU(2)$, so the nonabelian action is pseudoreal up to abelian twist. The first intrinsically complex case is therefore the triplet ($\chi = 3$), giving $SU(3)$.
- The minimal coupled carrier for both is $\mathbb{C}^3 \otimes \mathbb{C}^2$, giving coupled edge capacity $\chi_{\text{cpl}} = 6$.
- The block-diagonal faithful representation $\mathbb{C}^3 \oplus \mathbb{C}^2$ of $S(U(3) \times U(2))$ has dimension 5, but it is not coupled and therefore is not the MAR minimizer.
- The maximal compact subgroup of $U(6)$ acting on $\mathbb{C}^3 \otimes \mathbb{C}^2$ with commuting actions is $(SU(3) \times SU(2) \times U(1))/(\text{finite center})$.
- The commutant of $SU(3) \times SU(2)$ inside $U(6)$ is exactly $U(1)$, so any connected abelian factor acting nontrivially on the coupled carrier is necessarily a single $U(1)$, and no additional continuous factors appear without increasing χ_{cpl}.
- Product structure is not separately assumed: it follows from the tensor product structure of the minimal coupled carrier.

Combined with Proposition 6.6 (hypercharges quantized in sixths from the realized spectrum), this yields:

$$G_{\text{phys}} = \frac{SU(3) \times SU(2) \times U(1)}{Z_6}.$$

The full proof route is given in the remainder of Section 6.2 and in the dedicated gauge-derivation surface.

3.7.3 Refinement stability and unprotected relevant operators

The refinement-stability used later in the gauge branch is contained in the third OPH axiom. Because the same finite local constraint family is preserved across cutoffs, the realized UV states lie in one common finite-dimensional MaxEnt family rather than in a new coupling space at every refinement step. The selected stable branch of this family is therefore not an extra axiom beyond the MaxEnt branch itself.

Concretely, if

$$\omega_\ell(\lambda) = Z_\ell(\lambda)^{-1} \exp\left(- \sum_x \sum_a \lambda_a O_a(x) - \sum_b \mu_b Q_b \right),$$

then any refinement channel $\Phi_{\ell \to L}$ compatible with Axiom 3 acts on the realized branch by an induced finite-dimensional map

$$\Phi_{\ell \to L}(\omega_\ell(\lambda)) = \omega_L(R_{\ell \to L}(\lambda)).$$

The OPH "refinement-stable" branch is simply a trajectory or invariant subset of this finite-dimensional multiplier map. Accordingly, when later sections speak of a refinement-stable directed colimit of sectors, the refinement-stable qualifier means persistence along this MaxEnt branch itself; the remaining extra ingredients are transportability, symmetry, and the bosonic fiber-functor clauses.

Relevant operators that are neither symmetry-forbidden nor retained in the selected constraint family are precisely the directions that try to push the flow off that stable branch.

Lemma 6.7 (refinement-stable MaxEnt branch forbids unprotected relevant operators). Assume the local finite-constraint MaxEnt/refinement branch of the third OPH axiom. Let $\mathcal{O}$ be a gauge-invariant Lorentz-scalar relevant deformation in the emergent EFT sense ($\Delta < 4$ in $3+1$D), allowed by symmetry and absent from the retained constraint family. Then the selected branch can keep the coupling of $\mathcal{O}$ at zero only if symmetry forbids that direction or the constraint family explicitly retains it. Otherwise generic refinement induces a nonzero component along that direction and the flow leaves the selected branch. This is a branch-persistence statement, not a universal entropy-ordering theorem for arbitrary off-branch phases.

Proof sketch. Linearize the induced refinement map $R_{\ell \to L}$ on the finite-dimensional multiplier space around the selected branch. A relevant operator gives an unstable direction with scaling exponent $y > 0$. If $\mathcal{O}$ is not fixed by symmetry and is not part of the retained constraint family, generic UV mismatch produces a nonzero component along that direction. Because $y > 0$, repeated coarse-graining amplifies the component and pushes the flow off the selected branch unless one fine-tunes it away at every scale or protects it by symmetry or by the declared constraint family. QED.

Corollary 6.8 (chirality selector). A gauge-invariant Dirac mass term is a relevant scalar. If both chiralities exist in conjugate representations, the mass term is allowed and will be generated under refinement unless symmetry-forbidden or explicitly retained as a protected constraint. Therefore keeping light fermions on the selected refinement-stable branch requires chiral matter content or an explicit protecting mechanism. QED.

3.7.4 Generation number from CP violation and refinement stability (derived from MAR)

Anomaly cancellation is generation-by-generation, so it does not fix the number of generations. The admissibility conditions constrain the window; MAR selects the minimum.

Proposition 6.9 (The number of generations is N_g = 3). Under (i) intrinsic CP violation in the quark sector, (ii) UV-completability of SU(2)_L (asymptotic freedom at one loop),

(iii) MaxEnt/refinement stability selecting minimal viable spectrum, and (iv) the derived N_c = 3 from Theorem 6.14, the generation number is

$$N_g = 3.$$

Inputs.

1. **Intrinsic CP violation exists** in the quark sector (empirical fact; also, the framework treats "intrinsic CP violation" as a selector input).
2. **UV-completability proxy**: SU(2)_L is asymptotically free at one loop in the emergent EFT.
3. **MaxEnt + refinement stability** penalizes unnecessary unfixed flavor structure, selecting the minimal viable spectrum.
4. Use the derived N_c = 3 from Theorem 6.14.

Step 1: CP violation lower bound. The number of physical CP-violating phases in an N_g × N_g CKM matrix is:

$$\#(\text{CP phases}) = \frac{(N_g - 1)(N_g - 2)}{2}.$$

- For N_g = 1, 2: this is 0 → **no intrinsic CP violation possible**.
- For N_g = 3: this is 1 → **intrinsic CP violation possible**.

So intrinsic CP violation requires:

$$N_g \geq 3.$$

Step 2: SU(2) asymptotic freedom upper bound. The one-loop coefficient is:

$$b_{\text{SU}(2)} = \frac{22}{3} - \frac{1}{3} N_g (N_c + 1) - \frac{1}{6},$$

where the final $-1/6$ is the contribution of one complex Higgs doublet. Asymptotic freedom means $b_{\text{SU}(2)} > 0$, i.e.,

$$N_g (N_c + 1) < \frac{43}{2}.$$

With N_c = 3, we have N_c + 1 = 4, so:

$$4N_g < \frac{43}{2} \quad \Rightarrow \quad N_g \leq 5.$$

Combining: 3 ≤ N_g ≤ 5.

Step 3: MAR selection. Given the allowed window {3, 4, 5}, MAR (fourth component of the complexity vector $C(\mathfrak{S})$) selects the smallest viable choice:

$$N_g = 3.$$

QED.

Why this is convincing.

- It predicts a **single integer**.

- It uses **two admissibility conditions** (CP violation exists; weak sector is UV-completable) plus MAR's lexicographic minimality.
- It is not a fit to a continuous number.
- Under the stated gauge-selection hypotheses, this is a derived result, not conditional.

3.7.5 Hilbert-space formulation of gluing data

Let {P_i} be a good cover of the screen. For each patch, fix a representation

$$\pi_i : \mathcal{A}_i \to \mathcal{B}(\mathcal{H}_i).$$

For each overlap, choose a unitary intertwiner

$$U_{ij} : \mathcal{H}_j \to \mathcal{H}_i$$

such that for all $O \in \mathcal{A}_{ij}$,

$$\pi_i(O) = U_{ij}\pi_j(O)U_{ij}^{\dagger}.$$

Normalize U_ii = 1 and U_ji = U_ij†.

Lemma 6.10 (centrality on triple overlaps). On a triple overlap define

$$\Omega_{ijk} := U_{ij}U_{jk}U_{ki}.$$

For all $O \in \mathcal{A}_{ijk}$,

$$\Omega_{ijk}\pi_i(O) = \pi_i(O)\Omega_{ijk}.$$

Proof. Conjugation by U_ki sends π_i(O) to π_k(O), by U_jk to π_j(O), by U_ij back to π_i(O). Thus conjugation by Ω_ijk fixes π_i(O), so Ω_ijk commutes with π_i(O). QED.

Lemma 6.11 (gauge behavior). If Ũ_ij = V_i U_ij V_j† with V_i acting trivially on overlap observables, then

$$\tilde{\Omega}_{ijk} = V_i\Omega_{ijk}V_i^{\dagger}.$$

In particular, if Ω_ijk is central, its class is gauge invariant. QED.

3.7.6 Loop obstruction class (central defect)

Assume the defect is central and write

$$\varphi_{ij} := \mathrm{Ad}(U_{ij}) \quad \text{on } \mathcal{A}_{ij}.$$

This is the abelian truncation of the full 2-group obstruction in Section 3.4.

Then there exist central unitaries z_{ijk} such that

$$\varphi_{ij}\varphi_{jk}\varphi_{ki} = \mathrm{Ad}(z_{ijk}).$$

Theorem 6.12 (loop-coherent gluing iff vanishing obstruction). The family $\{z_{ijk}\}$ is a Čech 2-cocycle, and its class $[z]$ is gauge invariant. On any quadruple overlap P_{ijkl},

$$z_{jkl}z_{ikl}^{-1}z_{ijl}z_{ijk}^{-1} = 1.$$

A loop-coherent global gluing exists iff [z] = 0.

Proof. Compare two parenthesizations of ϕ_ij ϕ_jk ϕ_kl ϕ_li on a quadruple overlap to obtain the cocycle condition above. Gauge changes shift z by a coboundary. If [z]=0, rephase by a 1-cochain to eliminate defects and obtain path-independent transport. Conversely, loop-coherent gluing implies z_ijk = 1. QED.

3.7.7 EFT reduction to anomaly cancellation

Assume ExtEFT: a low-energy 3+1D chiral gauge theory exists with group G. Then the obstruction class $[z]$ is structurally analogous to the EFT 't Hooft anomaly class and is expected to map to it after a separate anomaly-descent construction. That map is not constructed in this manuscript, so $[z] = 0$ is used here as an internal gluing/transportability condition rather than a proved equivalence to EFT anomaly cancellation.

3.7.8 Hypercharge from anomaly freedom and Yukawas

Theorem 6.13 (Hypercharge from anomaly freedom and Yukawas). Assume gauge group SU(N_c) × SU(2) × U(1)_Y and one generation of left-handed Weyl fermions (Q, u^c, d^c, L, e^c), with a Higgs doublet H and Yukawa terms

$$QHu^c, \qquad QH^\dagger d^c, \qquad LH^\dagger e^c.$$

Then anomaly freedom and Yukawa invariance fix the hypercharges up to an overall normalization, yielding the Standard Model pattern for N_c = 3.

Proof. Yukawa invariance gives

$$Y_u = -(Y_Q + Y_H), \quad Y_d = -Y_Q + Y_H, \quad Y_e = -Y_L + Y_H.$$

Anomaly cancellation yields

$$\begin{aligned} SU(2)^2U(1): &\quad N_cY_Q + Y_L = 0, \\ \text{grav}^2U(1): &\quad 2N_cY_Q + N_cY_u + N_cY_d + 2Y_L + Y_e = 0. \end{aligned}$$

Solving gives

$$Y_L = -N_cY_Q, \quad Y_H = N_cY_Q, \quad Y_u = -(N_c+1)Y_Q, \quad Y_d = (N_c-1)Y_Q, \quad Y_e = 2N_cY_Q.$$

With these relations, SU(N_c)^{2}U(1) and U(1)^3 anomalies vanish automatically. Fixing the normalization by Q = T_3 + Y and Q(ν_L)=0 gives

$$Y_Q = \frac{1}{2N_c}.$$

For N_c = 3,

$$Y_Q = \frac{1}{6}, \quad Y_L = -\frac{1}{2}, \quad Y_e = 1, \quad Y_u = -\frac{2}{3}, \quad Y_d = \frac{1}{3}, \quad Y_H = \frac{1}{2}.$$

Without Yukawas, the cubic anomaly leaves two discrete branches (Y_u, Y_d exchange). Yukawa invariance selects the branch with a single Higgs doublet. QED.

Corollary 6.13a (Exact rational hypercharges). With the derived N_c = 3, the hypercharge assignments are uniquely fixed to exact rational values:

$$Y_Q = \tfrac{1}{6}, \quad Y_L = -\tfrac{1}{2}, \quad Y_u = -\tfrac{2}{3}, \quad Y_d = \tfrac{1}{3}, \quad Y_e = 1, \quad Y_H = \tfrac{1}{2}.$$

Why this is convincing.

- These are **exact rationals**, not approximate numbers.
- Their ratios are fixed by anomaly freedom + Yukawa invariance, and the absolute lattice is fixed by the standard normalization, with no continuous parameters to adjust.
- This high-precision set of numbers strongly constrains the realized matter package and matches observation exactly.

3.7.9 Witten anomaly and the number of colors

Theorem 6.14 (The number of colors is N_c = 3). Under the gauge structure SU(N_c) × SU(2)_L × U(1)_Y with one left-handed quark doublet Q per color and one left-handed lepton doublet L per generation, the global SU(2) anomaly (Witten, 1982) requires N_c to be odd. Under MAR (lexicographic minimization of $C(\mathfrak{S})$, where N_c is the third component), this yields:

$$N_c = 3.$$

Inputs.

1. Low-energy gauge group contains an SU(2)_L factor and an SU(N_c) color factor.
2. The matter content per generation includes:
    - one left-handed quark doublet Q which is an SU(2) doublet and carries color,
    - one left-handed lepton doublet L which is an SU(2) doublet and color singlet.
3. **Witten's global SU(2) anomaly constraint** (Witten, 1982): the number of left-handed SU(2) doublets must be even.
4. **MAR** (Selection Axiom): among allowed values, the third component of the complexity vector $C(\mathfrak{S})$ is minimized.

Proof. Count SU(2) doublets per generation:

- Quark doublets: N_c copies (one per color),
- Lepton doublets: 1 copy.

Total doublets per generation:

$$N_c + 1.$$

Witten anomaly cancellation requires this to be even:

$$N_c + 1 \equiv 0 \pmod 2 \quad \Rightarrow \quad N_c \text{ is odd.}$$

The Witten constraint alone allows N_c ∈ {1, 3, 5, 7, ...}. N_c = 1 fails admissibility (SU(1) is trivial, no complex nonabelian charge type). MAR (input 4) then selects N_c = 3 as the smallest nontrivial value. QED.

Status. The Witten anomaly derives N_c odd. The specific value N_c = 3 follows from MAR applied to the admissible class. Under the stated gauge-selection hypotheses, this is derived, not assumed.

Why this is convincing.

- It predicts a **single integer** given the minimality selector.
- The odd constraint is independent of continuous parameters, RG running, masses, or Yukawa values.
- It cannot be adjusted without changing the basic notion of electroweak doublets and color replication.

3.7.10 Bond-dimension gatekeeping

In tensor-network or code realizations, gauge actions act on edge factors of size χ, so emergent compact gauge groups embed in U(χ). This shows a capacity constraint: accommodating SU(3) color and SU(2) weak factors shows $\chi \geq 6$ in the minimal case, consistent with the MAR-derived gauge group.

3.7.11 Inevitability of photon and graviton

The model requires photons and gravitons.

Photon inevitability chain:

1. Gauge-as-gluing states that overlap identifications have redundancy forming a local groupoid.
2. Theorem 2.3 (edge-center completion) decomposes collar Hilbert spaces into sectors labeled by boundary gauge representations.
3. Theorem 6.1 (Tannaka/DR reconstruction) recovers a compact gauge group G from the fusion rules of these edge sectors.
4. Corollary 6.1 (conditional on DHR transportability) reconstructs a field algebra with G as a local gauge symmetry.
5. For the Standard Model, G includes U(1)_em after electroweak symmetry breaking.
6. A gauge boson is the quantum of the gauge field. Once U(1)_em emerges from overlap redundancy, its gauge field exists, and its quantum (the photon) must exist.

The photon is not postulated. It is forced by the axioms through the chain above. The photon mediates the correlations between charged excitations in different patches; it is how the U(1) redundancy structure propagates through the algebra net.

Graviton inevitability chain:

1. Theorem 4.2 (BW_{S^2}) shows that on the stated OPH Bisognano–Wichmann refinement branch, and under collar Markov locality plus the controlled refinement-limit premises, modular flow on caps becomes geometric conformal dilation.
2. Theorem 4.3 identifies the induced kinematic group as $\mathrm{Conf}^+(S^2) \cong \mathrm{PSL}(2,\mathbb{C}) \cong \mathrm{SO}^+(3,1)$, the Lorentz group.
3. Theorem 5.1 shows that, under the stated scaling-limit null-stress and reference-state premises, the condition δS_gen = 0 implies the rest-frame first-variation Einstein relation, with overlap consistency upgrading it to the semiclassical Einstein equations in the EFT regime.
4. The metric tensor emerges as the compression of modular flow data, and its dynamics are fixed by entanglement equilibrium.
5. A dynamical metric in a quantum theory requires a spin-2 quantum field. Its quantum (the graviton) must exist.

The graviton is not postulated. It is forced by the axioms through the chain above. Diffeomorphism invariance emerges because the bulk spacetime description is a compression of screen data; different coordinate descriptions are redundancies in how that compression is presented.

3.7.12 Quotient-protected charge quantization

Theorem 6.19 (Charge quantization and no fractional color singlets). If the global gauge group is

$$G_{\text{phys}} = \frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{Z_6},$$

as derived in Proposition 6.6, then

All color-singlet states have integer electric charge.

Equivalently: no stable isolated particles with charges like ±1/3 can exist as color singlets.

Proof. The Z$_6$ quotient identifies the center elements (e^{2πi/3}, -1, e^{iπ/3}) ∈ SU(3) × SU(2) × U(1) with the identity. For a color-singlet state ($\tau = 0$), the SU(3) factor acts trivially. The remaining identification requires the SU(2) × U(1) quantum numbers to satisfy

$$(-1)^{2j} \cdot e^{i\pi n/3} = 1,$$

where j is the SU(2) spin and n = 6Y is the integer hypercharge label. This gives n ≡ -6j (mod 6), i.e., n ≡ 0 (mod 6) for integer j and n ≡ 3 (mod 6) for half-integer j. Equivalently: Y is integer when j is integer, and Y is half-integer when j is half-integer.

After electroweak breaking, Q = T$_3$ + Y. For integer j, T$_3$ ∈ ℤ and Y ∈ ℤ, so Q ∈ ℤ. For half-integer j, T$_3$ ∈ ℤ + 1/2 and Y ∈ ℤ + 1/2, so Q = (half-integer) + (half-integer) ∈ ℤ. In both cases, Q ∈ ℤ. QED.

Experimental status: No fractionally charged color-singlet particles have been observed. Three independent high-precision bounds confirm this:

1. **Neutrality of matter** (PDG 2024): The proton-electron charge sum satisfies

$$|q_p + q_e|/e < 1 \times 10^{-21},$$

confirming charge quantization to 21 decimal places.

2. **Fractional charge searches in bulk matter**: Silicone oil drop experiments limit fractionally charged particle abundance to

$$(\text{fractionally charged particles})/\text{nucleon} \lesssim 10^{-22}.$$

3. **Collider searches** (CMS, PRL 134, 2025): Exclusions for stable particles with $q \in [e/3, 0.9e]$ up to masses $\sim$ 640 GeV (95% CL).

Existing searches are fully consistent with these structural consequences.

These are exact symmetry- or quotient-protected outputs: the empirical content is that violations remain excluded up to available experimental sensitivity.

3.7.13 Coupling extraction from edge-sector probabilities

The edge-center completion (Theorem 2.3) yields sector probabilities p_α on collar boundaries. These probabilities encode the renormalized gauge coupling through a heat-kernel/Laplacian weighting law.

Abelian case (Z_n). For a Z_n gauge theory, the edge sectors are labeled by charge q ∈ {0, 1, ..., n-1}. The correct "Casimir" eigenvalue is the Laplacian eigenvalue of the boundary random walk:

$$\lambda_q = 4 \sin^2\left(\frac{\pi q}{n}\right).$$

Note: only in the limit n → ∞ and q ≪ n does λ_q ≈ $(2\pi q/n)^2 \propto q^2$. For finite n, the exact form is essential.

The sector probabilities follow a heat-kernel law:

$$p_q \propto e^{-t(\mu)\lambda_q},$$

where t(μ) is the "modular time" parameter encoding the scale. The extraction formula is:

$$t(\mu) = -\frac{\log(p_q/p_0)}{\lambda_q}, \qquad g^2_{\mathrm{ent}}(\mu) = \frac{t(\mu)}{2\pi}.$$

Consistency requires that t extracted from different charges q agrees; this has been verified numerically (see Section 6.14).

Electric-center measurement. The edge sectors are measured using the *electric-center* prescription. For a region A and boundary vertex v ∈ ∂A, define the restricted star operator:

$$Q_v^{(A)} = \prod_{\ell \in \mathrm{star}(v) \cap A} X_\ell^{\pm 1},$$

where X_ℓ is the shift operator on link ℓ. The sector projectors are:

$$P_{v,q} = \frac{1}{n} \sum_{m=0}^{n-1} \omega^{-mq} \left(Q_v^{(A)}\right)^m, \qquad \omega = e^{2\pi i/n},$$

and the probabilities are p_{v,q} = ⟨P_{v,q}⟩. This electric-center operator, built from X's rather than Z's, correctly captures the boundary gauge charge/flux that labels entanglement edge sectors.

Non-abelian generalization. For SU(N) gauge theories, the edge sectors are labeled by irreducible representations with probabilities:

$$p_j \propto d_j \, e^{-t(\mu) C_2(j)},$$

where d_j is the dimension and C_2(j) the quadratic Casimir. Extraction:

$$t(\mu) = -\frac{\log(p_j/p_0)}{C_2(j)}, \qquad g^2_{\mathrm{ent}}(\mu) = \frac{t(\mu)}{2\pi}.$$

Theoretical derivation. The heat-kernel law can be derived from the axioms once the edge-sector generator is known to take the quasi-local local-Gibbs form derived in Theorem 2.6.

Theorem 6.20 (Heat-kernel law from MaxEnt + gauge structure). Under the OPH axioms, the fixed-cutoff overlap-gauge realization of Sections 2.3 and 3.2, and the quasi-local local-Gibbs form supplied by Theorem 2.6, the edge-sector probability distribution satisfies:

$$p_R = \frac{d_R \, e^{-t\lambda_R}}{\sum_{R'} d_{R'} \, e^{-t\lambda_{R'}}}$$

where λ_R is the Laplacian eigenvalue on the R-isotypic component and t is determined by the collar Gibbs parameter.

Proof.

Step 1 (Edge Hilbert space). From the fixed-cutoff overlap-gauge realization, the edge degrees of freedom at a boundary circle Σ = ∂C live in a Hilbert space transforming under the gauge group G. Microscopically this is a finite-dimensional quantum-link edge space; the effective representation-theoretic description in the refinement limit is modeled by $L^2(G)$. By the Peter–Weyl theorem [30]:

$$L^2(G) \cong \bigoplus_R V_R \otimes V_R^*$$

where V_R is the carrier space of irrep R.

Step 2 (Gauge invariance). The Gauss law constrains physical states. For an entanglement cut at Σ, the physical edge Hilbert space decomposes as 𝓗_{edge}^{phys} = ⊕_R W_R where W_R contains states with flux in representation R.

Step 3 (Natural Hamiltonian). From the local-Gibbs form, the MaxEnt generator restricted to edge modes takes the form H_{edge} = Σ_R h_R P_R where P_R is the projector onto the R-sector. The key claim is that h_R = λ_R.

Justification: For a compact simple factor, the group Laplacian Δ_G = -Σ_a (T^a)^2 is the unique bi-invariant second-order differential operator up to scale. For a product group $G = \prod_i G_i$, the most general bi-invariant second-order operator is a positive linear combination $\sum_i c_i \Delta_{G_i}$, with one coefficient per factor. Any other gauge-invariant local choice would require higher derivatives, violating locality. For finite groups, the Cayley graph Laplacian plays the same role: λ_R = |S| - (1/d_R) Σ_{s∈S} χ_R(s).

Step 4 (MaxEnt selection). MaxEnt selects the Gibbs state:

$$\rho_{\text{edge}} = \frac{1}{Z} e^{-tH_{\text{edge}}} = \frac{1}{Z} \sum_R e^{-t\lambda_R} P_R.$$

Step 5 (Sector probabilities). The probability of sector R is p_R = Tr(ρ_{edge} P_R). The effective dimension for entanglement is d_R (not d_R^2) because we trace over one side of the cut. This gives:

$$p_R = \frac{d_R \, e^{-t\lambda_R}}{Z}.$$

QED.

Why the entropy rank is d_R (instead of d_R^2). The full edge space has dimension d_R^2 in sector R (from $V_R \otimes V_R^*$). Entanglement entropy, however, measures correlations *across* the cut. After tracing over one side, the reduced density matrix has effective rank d_R. Mathematically: in the Markov normal form, the edge factor on one side contributes $\log d_R$ to the entropy.

Status. The derivation is complete up to the factorwise Laplacian choice just stated. The quasi-local local-Gibbs generator used here is derived from Theorem 2.6: if MaxEnt constraints are expectations of finitely many quasi-local operators, the entropy maximizer is automatically a Gibbs state with a quasi-local generator. What remains is the specific *Laplacian form* of that generator; this follows from gauge invariance plus the factorwise uniqueness of the bi-invariant second-order operator on each compact simple factor.

Normalization anchor: 2D Yang-Mills. The parameter t can be exactly matched to a conventional coupling in 2D Yang-Mills, where the physical Hamiltonian is literally the group Laplacian:

$$H = \frac{g^2}{2}\Delta_G, \qquad \Delta_G \chi_R = -C_2(R)\chi_R \quad \Rightarrow \quad E_R = \frac{g^2}{2}C_2(R).$$

Euclidean evolution for "time" A (the area of a cylinder in 2D YM) gives weight(R) ∝ exp(-A E_R) = exp(-g^2 A C_2(R)/2). Comparing with the heat-kernel expansion K_t(U) = Σ_R d_R χ_R(U) e^{-t C_2(R)} yields the exact identification:

$$t_{\text{phys}} = \frac{g^2 A}{2} \quad \text{(in 2D YM, no ambiguity).}$$

This shows that the Laplacian + MaxEnt → heat-kernel structure is exact in a solvable continuum Yang-Mills case. The coefficient in front of C_2 is fixed. In any regime where the edge theory reduces to an effective 2D YM with known "Euclidean thickness" A_eff:

$$g^2(\mu) = \frac{2}{A_{\text{eff}}(\mu)} \cdot \frac{\Delta_R(\mu)}{C_2(R)},$$

and the RHS must be R-independent. This R-independence is an internal precision consistency test; the formula itself is the normalization map that connects t to the conventional gauge coupling.

3.7.14 Numerical validation of the heat-kernel law

The heat-kernel/Laplacian weighting of edge sectors has been validated in explicit 2D Z_n gauge models on closed geometries.

Model. A 2×2 periodic lattice gauge theory (8 links) with Z_n link Hilbert spaces and Hamiltonian:

$$H = -K\sum_p \mathrm{Re}(B_p) - h\sum_\ell \mathrm{Re}(X_\ell) - \Gamma\sum_v \mathrm{Re}(A_v),$$

where X_ℓ is the Z_n shift on link ℓ, B_p is the oriented plaquette operator (product of Z's around plaquette p), and A_v is the oriented star/Gauss operator (outgoing X, incoming X†). With K = 1 and Γ = 5, the ground state satisfies ⟨A_v⟩ = 1 at all vertices to numerical precision.

Region and edge operator. Region A consists of links whose tail has x = 0 ("half-lattice" cut). At each boundary vertex v, the electric-center edge charge is the restricted star Q_v^{(A)} = ∏_{ℓ ∈ star(v) ∩ A} X_ℓ^{±1}.

Results for $\mathbf{Z_2}$. With $\lambda_1 = 4\sin^2(\pi/2) = 4$:

h	p_0	p_1	t	g_ent
0.5	0.8266	0.1734	0.391	0.249
1.0	0.9612	0.0388	0.803	0.357
2.0	0.9917	0.0083	1.194	0.436

Results for $\mathbf{Z_3}$ (overconstrained test). With $\lambda_1 = \lambda_2 = 4\sin^2(\pi/3) = 3$:

h	p_0	p_1	p_2	t(q=1)	t(q=2)	g_ent	m_plaq
0.2	0.4395	0.2803	0.2803	0.1500	0.1500	0.154	2.22
0.5	0.7509	0.1245	0.1245	0.5989	0.5989	0.309	1.75
1.0	0.9606	0.0197	0.0197	1.2956	1.2956	0.454	4.07

h	p_0	p_1	p_2	t(q=1)	t(q=2)	g_ent	m_plaq
1.5	0.9851	0.0074	0.0074	1.6288	1.6288	0.509	7.06
2.0	0.9921	0.0039	0.0039	1.8440	1.8440	0.542	10.10

The equality $p_1 = p_2$ is exact (charge conjugation symmetry in Z_3). The equality t_{q=1} = t_{q=2} is the crucial **overconstrained** check: at h = 1.0, extracting t from q = 1 and q = 2 independently gives t_{q=1} ≈ 1.2956389318579 and t_{q=2} ≈ 1.2956389318521. The agreement to ~10^{-14} (machine precision) confirms that the edge distribution genuinely follows the heat-kernel/Laplacian form.

Region-choice robustness. At h = 1, the extracted g_ent is nearly independent of region size:

- 2 links (one vertex's outgoing links): g_ent ≈ 0.453
- 4 links (half-lattice): g_ent ≈ 0.454
- 6 links (three vertices): g_ent ≈ 0.453

This locality confirms that the coupling is dominated by physics near the cut, not global bookkeeping, exactly what is expected if this behaves like a local QFT observable.

Results for Z_5 (golden ratio test). The Z_5 case provides a stringent test because the Laplacian eigenvalues have a distinctive ratio involving the golden ratio $\phi = (1+\sqrt{5})/2$:

$$\lambda_q = 4\sin^2\left(\frac{\pi q}{5}\right), \qquad \frac{\lambda_2}{\lambda_1} = \frac{\sin^2(72°)}{\sin^2(36°)} = \phi^2 \approx 2.618.$$

This ratio distinguishes the Laplacian law from naive alternatives: a linear model (λ_q ∝ q) would predict ratio 2, while a quadratic model (λ_q ∝ q^2) would predict ratio 4.

Simulations on a 2×2 torus in the dual/flux basis (125 states in the zero-winding sector) give:

h	Measured ratio ln(p_2/p_0)/ln(p_1/p_0)	Deviation from ϕ^2
0.5	2.25	14%
1.0	2.51	4%
2.0	2.619	< 0.1%

In the weak-field limit (h → 0, strong magnetic coupling), the simulation converges to the golden ratio squared. This confirms that the vacuum entanglement spectrum encodes the precise geometric structure of the gauge group Laplacian.

Significance. This validates the mathematical law (sector probabilities weighted by Laplacian eigenvalues) in explicit 2D gauge-invariant models with non-flat sector distributions. The Z_3 and Z_5 tests are structurally identical to SU(2)/SU(3): multiple irreps overconstrain the slope, and agreement confirms the mechanism works before jumping to nonabelian groups.

Results for S_3 (first nonabelian test). The abelian tests above use charge-sector projectors that reduce to Fourier modes. For nonabelian groups, the edge-sector projector must be generalized to character projectors:

$$P_{v,R} = \frac{d_R}{|G|}\sum_{h\in G}\chi_R(h^{-1})Q_v^{(A)}(h),$$

where d_R is the dimension of irrep R, χ_R is its character, and Q_v^{(A)}(h) is the restricted gauge action at boundary vertex v acting only on links in region A.

For S_3 (the smallest nonabelian group, order 6), there are three irreps: trivial (d=1), sign (d=1), and standard (d=2). The Cayley-graph Laplacian eigenvalues for the transposition generating set are:

$$\lambda_{\text{triv}} = 0, \qquad \lambda_{\text{sign}} = 6, \qquad \lambda_{\text{std}} = 3.$$

Exact reduction on one plaquette. For the single-plaquette model (4 links), imposing Gauss's law at all vertices means the physical wavefunction depends only on the plaquette holonomy's conjugacy class. Since S_3 has exactly 3 conjugacy classes, the gauge-invariant Hilbert space is 3-dimensional, spanned by the character states {|χ_R⟩}. In this basis, the edge-sector probabilities are exactly p_R = |c_R|^2 where $|\psi_0\rangle$ = Σ_R c_R |χ_R⟩. This is not an approximation; it is an exact identity for the one-plaquette gauge-invariant sector.

The heat-kernel ansatz predicts p_R ∝ d_R exp(-t λ_R). Extracting t independently from the sign and standard irreps provides an overconstrained test: the ratio λ_sign/λ_std = 6/3 = 2 is a parameter-free prediction. Results from a single-plaquette S_3 lattice gauge model (K=1, Γ=5):

h	p_triv	p_sign	p_std	t (sign)	t (std)	Δt/t	log-ratio
0.5	0.909	0.0013	0.089	1.09	1.01	8.4%	2.17
1.0	0.980	7.5×10^{-5}	0.020	1.58	1.54	2.8%	2.06
2.0	0.996	4.3×10^{-6}	0.004	2.06	2.04	1.0%	2.02
5.0	0.9993	1.0×10^{-7}	0.00066	2.68	2.67	0.3%	2.006
12	0.9999	3.0×10^{-9}	0.00011	3.27	3.27	0.1%	2.002
100	1.0000	6.1×10^{-13}	2.0×10^{-6}	4.69	4.69	0.009%	2.0002

The "Δt/t" column shows the fractional difference (t_sign - t_std) / $\bar{t}$. The "log-ratio" column shows log(p_sign/p_0) / log(p_std/(2 p_0)), which should equal λ_sign/λ_std = 2 if the heat-kernel form holds exactly.

As h increases, both diagnostics converge: Δt/t drops below 10^{-4} and the log-ratio approaches 2.000. This is exactly the expected behavior: finite-size corrections are largest at strong coupling; the heat-kernel form becomes exact as the perturbative regime is approached.

This is the first nonabelian validation of the edge-sector extraction mechanism. The structure (character projectors, Laplacian eigenvalues from the group's Cayley graph, overconstrained t extraction) is identical to what will be used for SU(2) and SU(3).

Parameter-free predictions for SU(2) and SU(3). The heat-kernel law yields exact, parameter-free ratio predictions that require no scheme matching. Define the "Casimir log-gap":

$$\Delta_R \equiv \ln\left(\frac{p_0}{d_0}\right) - \ln\left(\frac{p_R}{d_R}\right) = t\, C_2(R).$$

Ratios of Δ_R cancel all unknowns (t, partition function):

$$\frac{\Delta_{R_1}}{\Delta_{R_2}} = \frac{C_2(R_1)}{C_2(R_2)} \quad \text{(exact, parameter-free).}$$

SU(2) predictions. Irreps labeled by spin j have d_j = 2j+1 and C_2(j) = j(j+1). The framework predicts:

- $\Delta_1/\Delta_1/_2$ = 2/(3/4) = **8/3 ≈ 2.667**

- $\Delta_3/_2/\Delta_1/_2 = (15/4)/(3/4) = \mathbf{5}$
- $\Delta_3/_2/\Delta_1 = (15/4)/2 = \mathbf{15/8 = 1.875}$

SU(3) predictions. Irreps labeled by Dynkin indices (p,q) have $C_2(p,q) = (p^2 + q^2 + pq + 3p + 3q)/3$. Using the fundamental **3** = (1,0) with $C_2 = 4/3$ as the reference:

- $\Delta_8/\Delta_3 = 3/(4/3) = \mathbf{9/4 = 2.25}$
- $\Delta_6/\Delta_3 = (10/3)/(4/3) = \mathbf{5/2 = 2.5}$
- $\Delta_{10}/\Delta_3 = 6/(4/3) = \mathbf{9/2 = 4.5}$
- $\Delta_{15}/\Delta_3 = (16/3)/(4/3) = \mathbf{4}$
- $\Delta_{27}/\Delta_3 = 8/(4/3) = \mathbf{6}$

These are the SU(2)/SU(3) analogs of the Z_5 golden-ratio test: exact rational numbers fixed entirely by group theory, with no adjustable parameters.

Preliminary SU(3) results. A one-plaquette SU(3) "quantum link" model (finite truncated irrep basis, n_max = 12, $\kappa = 2$) has been used to extract t from 14 different irreps simultaneously. The results show internal consistency at the 1-3% level:

bare g^2	extracted t (mean±std)	g_ent	gap
0.3	0.314±0.0005	0.224	1.92
0.5	0.539±0.0025	0.293	1.83
0.8	0.896±0.012	0.378	1.72
1.0	1.144±0.025	0.427	1.64

The standard deviation across irreps provides a built-in error estimate. This is a QCD proton-physics surrogate rather than a full proton calculation; it lacks dynamical quarks and operates on a single plaquette. It nevertheless demonstrates that the nonabelian extraction machinery produces self-consistent outputs without tuning.

Extracting the normalization factor A_eff. The 2D YM anchor (Section 6.13) gives $t = g^2 A / 2$, so the "effective Euclidean thickness" is

$$A_{\text{eff}} = \frac{2t}{g^2}.$$

Computing this from the SU(3) table:

bare g^2	extracted t	A_eff
0.3	0.314	2.093
0.5	0.539	2.156
0.8	0.896	2.240
1.0	1.144	2.288

Mean: A_eff ≈ 2.19 with point-to-point scatter ~4%.

Extrapolation to weak coupling. The systematic drift in A_eff shows fitting $A\_eff(g^2) = A_0 + a \cdot g^2$. A weighted linear fit gives:

$$A_0 = 2.004 \pm 0.012$$

with χ^2/dof ≈ 0.09, indicating excellent consistency. This strongly shows that, in this toy UV completion, the "missing normalization" converges to A_eff → 2 as $g^2 \to 0$.

The normalization factor behaves like a quasi-constant and extrapolates to a simple value (≈ 2) in the weak-coupling limit. This provides a concrete path to absolute coupling predictions: once A_eff is determined from microphysics, the conversion g^2 = 2t/A_eff fixes the gauge coupling without additional free parameters.

Internal validation summary. The heat-kernel law has been validated with increasing precision across multiple gauge groups:

- $\mathbf{Z_3}$: Overconstrained t extraction (q=1 vs q=2), precision ~10^{-14}
- $\mathbf{Z_5}$: Golden ratio squared ($\lambda_2/\lambda_1 = \phi^2$), precision 0.04%
- $\mathbf{S_3}$: Casimir log-ratio (λ_sign/λ_std = 2), precision 0.01%
- **SU(3)**: 14-irrep simultaneous extraction, precision 1-3%

The Z_3 test achieves machine precision because it is exactly overconstrained. The Z_5 and S_3 tests converge to their predicted ratios as coupling decreases. This provides strong internal validation of the mechanism "MaxEnt + Laplacian => heat-kernel sector weights" before applying it to physical gauge groups.

3.7.15 IBM Quantum Cloud hardware benchmarks

The lattice calculations above are internal numerical validations. The IBM Quantum Cloud benchmarks add a hardware check of whether OPH-motivated reduced-sector structures survive preparation and readout on real superconducting-qubit devices. These runs are small-sector consistency benchmarks rather than a standalone confirmation of the full framework. Their value is that they probe the same recoverability and heat-kernel observables on physical hardware and in simulation.

Detailed circuits, representative raw outputs, and rerun instructions are public in the project write-up `extra/IBM_QUANTUM_CLOUD.md` and the associated code/data bundle `code/ibm_quantum_cloud/`.

- **Stage 1 (Markov/recoverability benchmark).** On `ibm_marrakesh` and `ibm_fez`, the structured state reconstructs below both controls in conditional mutual information and above both controls in Petz fidelity. On `ibm_marrakesh`, the observed ordering is $0.2309 < 0.3890 < 0.9474$ for CMI and $0.9297 > 0.8649 > 0.6066$ for Petz fidelity; on `ibm_fez`, it is $0.1498 < 0.4992 < 0.9166$ and $0.9479 > 0.8117 > 0.6449$. This is the OPH-favored direction on two independent real backends.
- $\mathbf{Z_3}$ **hardware sanity check.** The overconstrained Z_3 extraction passes cleanly on hardware: across prepared $t = 0.30, 0.60, 0.90$, the mean extracted t remains within about 0.02 of the target, the two independent extractions agree to within 0.0142, 0.0034, and 0.0043, and leakage stays below 0.1%. This shows that the reduced-sector preparation and readout path is internally coherent on-device before sharper ratio claims are interpreted.
- $\mathbf{Z_5}$ **exact-ratio test.** The parameter-free OPH target is $\Delta_2/\Delta_1 = \varphi^2 \approx 2.618$. Across repeated sweeps on `ibm_marrakesh` and one replication on `ibm_fez`, the best `ibm_marrakesh` points land within 0.8%, 1.2%, and 1.5% of the target, while a focused high-shot rerun at $t = 0.90$ lands within 2.6%. The `ibm_fez` replication is noisier but remains in the same neighborhood rather than producing a clean contradictory ratio.
- $\mathbf{S_3}$ **nonabelian ratio test.** The first nonabelian hardware run revealed a real layout-dependent bias. After reversing the qubit layout, the mitigated ratio moves from 1.8724 to 2.0299 against the OPH target 2.0000, and a direct confirmation run returns 2.0657. The

nonabelian target reappears once the hardware mapping is corrected, indicating an identifiable device/layout effect instead of a structural failure of the OPH prediction.

These IBM runs are best read as hardware benchmarks for the local OPH edge-sector picture. They show that the OPH-favored recoverability ordering survives on two real backends, the abelian exact-ratio tests land near the predicted parameter-free values, and the first nonabelian test returns to the expected ratio once a diagnosed layout bias is corrected. They add real-device evidence to the simulation-based checks, but they do not validate the full framework by themselves.

3.8 Status, Tests, and Open Problems

3.8.1 Classification of results and remaining dependence

This section states which node of the program has been proved, which standard mathematics is imported, and which extra physical premises or external inputs remain load-bearing. The documentary node labels D1–D12 match the ledger in Section 3.2.6. The three-phase boundary is

$$\text{Phase I} = (D1\text{–}D5) \cup (D7\text{–}D9), \qquad \text{Phase II} = D6 \cup D10, \qquad \text{Phase III} = D12.$$

Node	Representative outputs in this manuscript	Standard mathematics used	Extra physical premises / external inputs required	Claim tier
D1	Theorem 3.1 schedule-independent normal form	Newman's lemma	finite patch net, termination, and local confluence	Phase I structural theorem
D2	EC decomposition, exact/approximate Markov collar, and generalized-entropy split	HJPW; Petz; Fawzi–Renner	explicit recoverability control when exact identities are used; coarse-grained $A/(4G)$ dictionary	Phase I structural lemma/proposition layer
D3	Theorems 4.2–4.3 Lorentz branch	Bisognano–Wichmann modular geometry; conformal-group classification	rotationally invariant constraints together with T1 and T3	Phase I conditional scaling-limit theorem
D4	Section 5.2 null modular bridge	Borchers–Wiesbrock; distribution theory for half-line differentiation	the theorem-local inherited-strip decomposition and exact-or-controlled Markov hypotheses of Section 5.2, plus the downstream density-upgrade and relativistic null-stress-identification premises; the endpoint/propagation controls are internal to the local finite-constraint MaxEnt branch	Phase I conditional bridge theorem
D5	Theorem 5.1 rest-frame Einstein relation plus the conditional tensor upgrade of Corollary 5.1	Jacobson small-ball mathematics; null-to-tensor reconstruction modulo a metric term	fixed-cap stationarity, the locally Lorentzian $d = 4$ scaling regime, and the all-directions/all-reference-states premise used in the tensor upgrade	Phase I conditional scaling-limit theorem/corollary

Node	Representative outputs in this manuscript	Standard mathematics used	Extra physical premises / external inputs required	Claim tier
D6	capacity/Λ statements and the separate capacity-based neutrino estimate	de Sitter entropy relation and dimensional analysis	external input N_{scr} and the screen-capacity identification	Phase II input-dependent corollary / estimate
D7	Theorem 6.1 compact gauge reconstruction	Doplicher–Roberts / Tannaka reconstruction	bosonic symmetric transportable colimit and fiber-functor premises; the refinement-stable qualifier is internal to the local finite-constraint MaxEnt/refinement branch	Phase I conditional structural theorem
D8	product gauge structure up to finite quotient under MAR	compact Lie representation classification; Schur's lemma	MAR, connected admissible class, one connected abelian factor, and vanishing relevant transport obstruction	Phase I realized-branch theorem
D9	realized $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)/\mathbb{Z}_6$, exact hypercharges, the counting chain $N_g = 3$ then $N_c = 3$, and no gauge-mediated proton decay as a product-group corollary	anomaly algebra; Witten anomaly; CKM CP counting	realized one-generation/one-Higgs package and MAR admissibility premises	Phase I realized-branch theorem/corollary chain
D10	gauge-coupling numerical closure and target-free electroweak repair theorem of the D10 calibration branch	RG running and matching conventions	external input P, pixel closure, and the printed threshold conventions	Phase II calibration sector
D12	Koide/charged leptons, textures, dark-sector, baryogenesis, spectroscopy, string/worldsheet, and other downstream continuations	branch-specific EFT and phenomenological manipulations	explicit additional ansätze beyond D6, D9, or D10	Phase III phenomenological continuation / program branch

The ledger implies the following reading rules.

1. Phase I, the recovered core carried by this manuscript, is D1–D5 together with D7–D9. These are the structural or realized-branch claims defended at theorem/corollary level once all stated premises are included.

2. Phase II contains the separate input-dependent capacity corollary D6 and the D10 calibration sector. Failure there does not by itself erase Phase I.

3. Phase III contains D12 continuations only. Nothing in D12 is theorem-level unless its extra ansätze are replaced by proved premises.

4. **Existence / strange-loop narratives**: self-simulation or strange-loop closure stories are interpretive epilogues only. The internally available ingredient is a theorem-usable OPH state-and-law habitat built from overlap-consistent patch states together with finite Axiom-3 law coordinates. The recovered core does *not* derive an existence theorem, an OPH-internal closure map on an invariant admissible observer-supporting sector, or any uniqueness/stability result for such a map.

The corresponding low-energy effective-action form is

$$\mathcal{L}_{\mathrm{eff}}^{\mathrm{OPH}} \approx \sqrt{-g}\left[\frac{1}{16\pi G}(R - 2\Lambda) + \mathcal{L}_{\mathrm{SM}}^{\text{realized branch}}\right] + \sum_i \frac{c_i}{M_*^{\Delta_i - 4}}\mathcal{O}_i.$$

Here $\mathcal{L}_{\mathrm{SM}}^{\text{realized branch}}$ denotes the realized $\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)/\mathbb{Z}_6$ sector with the exact hypercharge lattice and the realized counting chain $N_g = 3$ then $N_c = 3$. Phase I secures the conditional Einstein branch and the realized Standard Model structural chain, D6 supplies Λ once the screen-capacity identification is adopted, and the higher-dimension operators absorb UV-sensitive or downstream corrections not fixed by the recovered core itself.

Symmetry-protected exact zeros such as the masslessness of the photon, gluons, and graviton, together with the absence of gauge-mediated proton decay, are corollaries inside Phase I. They do not define separate axiom-only nodes.

3.8.2 Structural assessment

- **Dynamics**: The GR chain requires modular covariance plus the fixed-cutoff null-surface modular bridge. Section 5.2 carries that bridge through the derived positive null-translation stage on the OPH geometric scaling branch. The remaining inputs are modular covariance, the downstream density upgrade, and the relativistic null-stress identification.
- **Gauge structure**: The gauge group is reconstructed from sector fusion, and the anomaly/gluing link is precise. Axiom 5 (MAR) fixes the SM factors; the CP/UV admissibility window and MAR give $N_g = 3$, and Witten's parity constraint with that realized generation count gives $N_c = 3$. Proposition 6.1a gives DHR transportability $\Leftrightarrow$ $[z] = 0$ on the central branch, while Corollary 3.4a gives the higher-gauge analogue $q_\Sigma = 0$ on the genuinely noncentral branch. The chirality selector still has to be justified in explicit models.
- **Microscopic theory**: Quantum link models (Section 2.6) provide explicit fixed-cutoff realized presentations and give EC + Markov collars automatically; they are not proofs of a unique microscopic completion. Open items are (i) a microscopic derivation of Recoverable Generalized Entropy, and (ii) a continuum limit in which modular flow becomes geometric, namely the **T1+T3** branch.
- **Loop gluing beyond central defect**: At fixed cutoff the genuinely noncentral branch is closed by the crossed-module higher-gauge EC and transport package; what remains open is the refinement-limit transportable-sector lift together with quantitative matching to EFT anomalies.

3.8.3 Testable Items and Conditional Phenomenological Branches

The framework makes several directly testable statements, together with weaker continuation-level templates that can also be confronted with data:

GW horizon spectroscopy comb (Section 5.11; continuation-level). If the extra discrete-horizon premises of Section 5.11 hold, the log-integer area spectrum gives discrete resonant frequencies for Kerr black hole horizons:

f_{k,m}(M,χ) = (m Ω_H)/(2π) + (c^3 g(χ))/(16π^2 GM) · ln(k) for k = 2, 3, 4, ...

After rescaling by remnant parameters, all events should stack at universal coordinates x_k = ln(k)/(8π). This is checkable with public LIGO/Virgo data as a continuation-level test of that extra branch. Absence of coherent stacking would challenge the discrete-horizon continuation, not the recovered-core package by itself.

Discrete Hawking comb (Section 5.11; continuation-level). If the same discrete-horizon branch is realized, primordial black holes in the final evaporation stage should show comb structure at E_k/E$_2$ = ln(k)/ln(2). Current PBH burst searches (Fermi, H.E.S.S.) can constrain this with dedicated template analysis.

Casimir ratio precision (Section 8.1). Future lattice measurements of SU(3) edge-sector probabilities should confirm $\Delta_8/\Delta_3 = 9/4$ exactly, not 2.67 (dimension-only) or 5.06 (Casimir-squared). The full set of parameter-free SU(3) ratio predictions is:

- $\Delta_8/\Delta_3 = 9/4 = 2.25$
- $\Delta_6/\Delta_3 = 5/2 = 2.5$
- $\Delta_{10}/\Delta_3 = 9/2 = 4.5$
- $\Delta_{15}/\Delta_3 = 4$
- $\Delta_{27}/\Delta_3 = 6$

These exact rationals are fixed entirely by group theory (Casimir eigenvalue ratios), with no adjustable parameters. Any deviation identifies a measurement contradiction with the heat-kernel edge-sector mechanism.

Conditional $\mathbf{Z}_6$ center-label entropy scale (Section 6.18). Under the uniform sixfold center-label ensemble used in the flavor continuation, the associated entropy scale is $\log_2 6 \approx 2.585$ bits. This is a phenomenological ansatz tied to the quotient label set, not a theorem-level direct observable of the quotient alone.

Black hole spectroscopy secondary structure (Section 5.11). Beyond the headline log-integer comb, the framework predicts rigid secondary structure:

1. *Universal energy ratios*: E_k/E_2 = ln(k)/ln(2) exactly. For example, E_3/E_2 = ln(3)/ln(2) ≈ 1.585 is parameter-free. This arithmetic pattern of ratios distinguishes OPH from other "quantized area" proposals that have different functional forms or free spacing parameters.

2. *Mass-independent fractional linewidth*: The intrinsic linewidth $\Gamma/\Delta E_k \approx$ 3–5% is approximately independent of black hole mass. This is a sharp shape prediction constraining line positions and line profiles.

3. *Fixed weight hierarchy*: Line weights follow a (k-1)/k pattern from detailed balance in the log-integer transition rule, on top of the GR greybody envelope. High-k lines asymptote in strength in a specific, counting-driven way.

Inequality bounds on GR deviations (Section 5.8). The modular additivity defect satisfies the exact identity ⟨ΔK⟩ = -I(A:D|B), where I(A:D|B) is the conditional mutual information. Under the Markov/mixing assumptions, this defect is exponentially small in collar thickness:

|⟨ΔK⟩| ≤ 2|A| · η^{w/ξ}

This propagates into an explicit upper bound on how far the Einstein equation can deviate from GR in regimes where the emergence proof applies. Unlike typical beyond-GR frameworks that postulate corrections, OPH provides a quantitative ceiling: given the information-theoretic primitives, corrections decay exponentially with collar width. This "UV ignorance → rigorous inequality" structure is distinctive.

Shared excitation dictionary benchmark (Section 6.21). The same-family continuations are organized by one common excitation dictionary instead of a primitive one-number suppression law. The proof-facing data are the same-label gap/overlap scalars

$$(g_e, \omega_e), \qquad d_e = 1 - \omega_e, \qquad q_e = \sqrt{g_e d_e}, \qquad \eta_e = \log q_e - \frac{1}{3}\sum_f \log q_f,$$

equivalently the centered eta-class $[\eta_e]$ or the normalized weights

$$\mu_e := \frac{e^{\eta_e}}{\frac{1}{3}\sum_f e^{\eta_f}}.$$

Let s_e denote the realized sector-charge step carried by the arrow e. The internal integer data are the excitation multiplicities

$$m_{\gamma,e} \in \mathbb{Z}_{\geq 0}$$

with which a given readout path γ traverses the realized same-label arrows, equivalently the integers satisfying

$$s_\gamma = \sum_e m_{\gamma,e} s_e,$$

and the corresponding excitation action

$$A_\gamma = -\sum_e m_{\gamma,e} \log q_e.$$

Old “base-6 exponents” are only the compression

$$n_\gamma^{(6)} := A_\gamma / \log 6,$$

which becomes an actual defect count only after the extra sixfold-uniform collapse compatible with the realized $\mathbb{Z}_6$ quotient. The intrinsic neutrino lane factors directly through the same-label scalar certificate, the quark mass-side continuation compresses to $(\Delta_{ud}^{\text{overlap}}, \eta_Q^{\text{centered}})$ and the emitted ray $\mathcal{R}_{D12}^{ud}$, and the charged lane fixes centered logs once a charged source pair exists. The remaining blockers therefore sit above the dictionary: $\widehat{C}_e^{\text{cand}} \to \widehat{C}_e$ and $\mu_{\text{phys}}(Y_e)$ for charged leptons, the intrinsic scale law on the emitted quark ray, and the neutrino absolute-scale correction scalar C_ν. The scalar benchmark $\varepsilon = 1/6$ survives only as the extra uniform sixfold collapse of this richer dictionary, so exponents $-\log(y_f)/\log 6$ are compare-only summaries of excitation action, not recovered-core outputs or independent flavor inputs.

No gauge-mediated proton decay (Section 6.11). If the gauge group is genuinely a product (from sector factorization), coupling unification is geometric (shared edge diffusion parameter) rather than simple-group embedding. This predicts unification-like coupling relations *without* GUT leptoquark bosons, hence no gauge-mediated proton decay. The combination "coupling unification + no proton decay" is a crisp discriminator against classic GUT predictions.

3.8.4 What is not predicted (gaps)

Section 3.8.1 fixes the three-phase boundary. The list below collects the items that remain outside Phase I even when they are discussed elsewhere in the manuscript.

Not predicted by the recovered core.

- **Flavor and Yukawa data**: Yukawa matrices, flavor hierarchies, charged-lepton/Koide fits, quark-texture exponents, CKM/PMNS closure, and related downstream flavor fits are not Phase-I outputs. The common proof-facing family base is the same-label excitation dictionary

  $$(g_e, \omega_e) \longmapsto d_e = 1 - \omega_e, \qquad q_e = \sqrt{g_e d_e}, \qquad [\eta_e].$$

  Its internal integer data are the realized traversal multiplicities

  $$m_{\gamma,e} \in \mathbb{Z}_{\geq 0}$$

with

$$s_\gamma = \sum_e m_{\gamma,e} s_e.$$

The associated excitation action is

$$A_\gamma = -\sum_e m_{\gamma,e} \log q_e,$$

from which any compare-only base-6 compression is formed. What remains open is the lane-specific readout closure above that dictionary: on the emitted local quark surface the same-label selector is fixed to σ_{ref}, which leaves the physical CKM-shell no-go intact on the selected D12 sheet, and an explicit same-family underdetermination theorem says that the mass side emits only the one-parameter ray $D12^{\mathrm{mass}}_{ud}$. The next exact quark object is the intrinsic scale law on the emitted D12 mass ray. The charged-lepton lane waits on promotion of $\widehat{C}^{\mathrm{cand}}_e$ and then one descended physical affine scalar $\mu_{\mathrm{phys}}(Y_e)$, from which the uncentered lift, determinant-line section, and affine anchor A_{ch} follow canonically. The neutrino continuation branch waits on the reduced bridge-correction invariant C_ν above the emitted proxy, with $B_\nu := \lambda_\nu q^{p_\nu}_{\mathrm{mean}}/m_{\star,\mathrm{eV}}$ retained as the paper-facing amplitude parameterization. The $\varepsilon = 1/6$ texture language is the extra uniform-center-label collapse of this richer dictionary and is not a recovered-core theorem.

- **Hadron masses and resonances**: the particle branch supplies the seeded-family, correlator, and stable-channel readout bridge, but promotable hadron rows remain execution-contract-frozen. They require one production backend export bundle, then real nonperturbative production computation and production systematics instead of an additional symbolic theorem alone, so they are not derived outputs of the particle-spectrum paper.

- **Neutrino data**: beyond the input-dependent, order-of-magnitude capacity estimate tied to N_{scr}, the recovered core does not derive neutrino masses or mixings. A continuation branch emits intrinsic masses, splittings, and PMNS data; its weighted-cycle branch reaches the physical PMNS/hierarchy regime, the normalized same-label overlap-defect weight section closes beneath that branch, and the absolute spectrum remains one reduced bridge-correction invariant C_ν short of theorem-grade closure above the emitted proxy, with the paper-facing amplitude $B_\nu := \lambda_\nu q^{p_\nu}_{\mathrm{mean}}/m_{\star,\mathrm{eV}}$ and the positive normalization scalar λ_ν induced only after that reduced bridge law is fixed. Because the positive selector segment is explicit, a compare-only two-parameter continuation adapter hits the representative central splittings exactly on that same family, but it stays strictly non-promotable while C_ν remains open. Together with the exact W/Z sidecar pair, the exact Higgs/top sidecar, and the exact charged/quark same-family witnesses, the exact non-hadron output lane is complete even though the theorem-grade charged, quark, and neutrino lanes remain open. The attached weighted-cycle stack factors through $q_e = q_{\mathrm{mean}} qbar_e$, so the manuscript does not derive a $qbar_e$-only bridge collapse law. The best closed constructive local object beneath that bridge is the defect-weighted same-label edge family $q_e = \sqrt{g_e d_e}$ together with its induced μ_e family, but that family does not emit C_ν or B_ν by itself.

- **Dark-sector phenomenology**: the sign, closure, and galaxy-scale response law of any modular-anomaly stress tensor are not derived by the recovered gravity chain.

- **Baryogenesis**: suppression-counting estimates are not a substitute for a derived out-of-equilibrium mechanism.

- **Black-hole spectroscopy**: any logarithmic line template requires an additional discrete-horizon branch.
- **String/worldsheet reorganization**: the Gross-Taylor-type interpretation [29] requires a controlled large-N_{edge} regime distinct from the proved compact-gauge chain, and any further lift to critical superstring structure remains a conjectural extension requiring extra worldsheet ingredients.
- **Other downstream continuations**: strong-CP proposals, proton-spin fractions, and similar late-stage continuations appearing on other OPH surfaces are not derived by the recovered core.

What remains open at the core/implementation boundary.

- **Scaling-limit control**: the Lorentz/null-modular/Einstein chain depends on the explicit scaling-limit premises declared earlier in the manuscript; proving or constructing that regime is an open task.
- **Fermionic gauge branch**: the manuscript proves compact gauge reconstruction only in the bosonic internal-gauge branch; a full fermionic/super-Tannakian extension remains open.
- **External inputs**: P and N_{scr} are not derived in this manuscript. The numerical value of Λ is not a local recovered prediction because local null data determine the Einstein equation only up to a metric term ϕg_{ab}; fixing Λ requires the global capacity input N_{scr}.
- **Calibration sector**: the D10 numerical implementation depends on supplement-backed running, matching, and threshold conventions. The displayed carrier closes its own exact W/Z chart, and the promoted target-free source-only repair theorem emits the public electroweak quintet on the Phase II calibration branch. The freeze-once coherent repair surface is retained only as compare-only validation; on that frozen authoritative surface the canonical W/Z pair is hit exactly, but it remains below the target-free theorem.
- **Hadron execution boundary**: the hadron lane is not blocked by one small symbolic gap. The next object is one production backend export bundle with publication-complete manifest provenance and real stable-channel correlator arrays; the normalized production dump and then production statistical/systematic control only follow after that bundle lands, so hadron rows remain outside the derived output surface.

A compact Higgs/top forward seed exists downstream of the D10 gauge core, but it remains a secondary quantitative descendant rather than a recovered-core theorem or a completed particle-spectrum closure. A separate compare-only inverse slice on the same D11 Jacobian can be solved to hit the canonical Higgs/top reference pair exactly, but that sidecar is not the live forward branch and does not change the theorem status of the public D11 rows.

Failures localize by tier. A failure in a Phase-III flavor, dark-sector, baryogenesis, spectroscopy, string, strong-CP, proton-spin, or similar continuation retracts that continuation only. A failure in D10 challenges the corresponding calibration sector. Only a failure in Phase I would overturn the recovered-core claim set.

3.8.5 Comparison with other unification approaches

Unified models attempting to tie together QFT, gravity, and SM structure tend to encounter a repeatable set of conceptual difficulties. This subsection examines how the observer-patch holography framework addresses these common pitfalls.

1. Subsystem factorization in gauge theory and gravity.

In gauge theories and gravity, the Hilbert space does not cleanly split as "inside $\otimes$ outside" across a cut. This infects entanglement entropy definitions, area terms, edge modes, and observable identification. Many unification attempts handwave this or patch it with conventions.

OPH route: The framework builds from a net of von Neumann algebras on patches plus overlap consistency. It does not start from naïve tensor factorization. The gauge-as-gluing + regulator package yields edge-center completion: a canonical block decomposition on collars where the center captures superselection data at the cut, and the state becomes (exactly or approximately) Markov across the collar. The entropy split S(ρ_C) = S_bulk + $\langle$L_C$\rangle$ follows from having a center with sector labels. This replaces the ad hoc "add an area term" move.

2. Modular Hamiltonian nonlocality.

Many entanglement-based gravity derivations depend on modular Hamiltonians that look like local stress-tensor charges (true only in special states/regions). In generic QFT states, modular Hamiltonians are nonlocal, making "first law of entanglement $\Rightarrow$ Einstein equation" arguments fragile.

OPH route: The Markov collar condition does heavy lifting: approximate Markov implies approximate modular additivity, with the defect controlled by conditional mutual information. This makes "modular locality" a controlled approximation. It is not treated as an assumption. On the stated Bisognano–Wichmann geometric branch of Theorem 4.2, the controlled tangent-half-space comparison then locks modular flow to geometric dilations with rigid 2π normalization.

3. Lorentz invariance as a derived output.

Discrete microscopic models generally break Lorentz symmetry, and many unified proposals simply postulate Lorentz invariance in the IR.

OPH route: Lorentz kinematics are tied to geometric modular flow on caps. On the stated BW geometric branch, once modular flow acts as conformal transformations on S^2, we get $\mathrm{Conf}^+(S^2) \cong \mathrm{PSL}(2,\mathbb{C}) \cong \mathrm{SO}^+(3,1)$, so the Lorentz group appears as a theorem-level output of modular structure on that branch. No external spacetime symmetry axiom is added.

4. Dynamics beyond kinematics.

Many approaches produce emergent geometry/kinematics but stall at dynamics: why Einstein's equations (with the right coefficient) rather than some other geometric PDE?

OPH route: The framework combines MaxEnt entanglement equilibrium, the derived K_C = 2πB_C structure with rigid normalization, and an EFT bridge identifying modular energy with stress-tensor charges. The null modular additivity route reduces much of that bridge to Markov/edge-center mechanisms on null strips plus explicit scaling-limit regularity premises. It does not rely only on "assume a UV CFT."

5. Gauge symmetry origin and compactness.

Most unification stories pick a gauge group and work out consequences. Emergent-gauge approaches sometimes produce noncompact groups or uncontrolled redundancies.

OPH route: Gauge symmetry is recast as redundancy in overlap identifications (gauge-as-gluing). From edge sectors and fusion, a tensor category is reconstructed; Tannaka-Krein / Doplicher-Roberts reconstruction then yields a compact group G given the categorical hypotheses. "Gauge symmetry" names the gluing redundancy at the conceptual level. "Compact group" is the mathematical form compatible with finite-dimensional sector/fiber-functor structure.

6. Massless photon and graviton usually hand-imposed.

Getting massless gauge bosons is easy if exact gauge invariance is assumed, but that restates the problem. Massless graviton is more delicate (mass terms, vDVZ discontinuity, strong coupling scales).

OPH route: Once gauge and diffeomorphism invariance are emergent redundancies of description (from gluing consistency / emergent geometry), hard mass terms are forbidden: "a coordinate system's Jacobian can't show up as a physical mass." These symmetry-protected zeros emerge from the same consistency machinery that gives the symmetries.

7. Global consistency, anomalies, and loop patching.

Building physics from local patches hits loop/holonomy problems: consistent gluing on a tree but obstructions around loops. These obstructions are often anomalies or global topological constraints.

OPH route: This is elevated to a first-class organizing principle: gluing data on overlaps defines cocycles; central defects define a Čech obstruction class [z] (and more generally a 2-group/crossed-module cocycle for noncentral defects). "Global consistency exists iff the obstruction class vanishes" becomes the universal statement. Anomalies become "failure to glue" rather than a mysterious quantum pathology.

8. Charge quantization without a GUT.

Without embedding into a simple GUT group, explaining charge quantization (why all isolated color singlets are integer charged) is awkward. Standard lore requires grand unification or monopoles.

OPH route: The framework leans on global group structure (the Z_6 quotient) and derives congruence/selection rules for allowed representations/hypercharges. This gives a structural explanation for integer-charged color singlets without introducing the proton-decay channel of simple-GUT models.

9. Coupling unification usually forces proton decay.

Traditional simple-group unification introduces leptoquark gauge bosons (X, Y) mediating proton decay. Experiment keeps pushing limits up, pressuring minimal GUTs.

OPH route: "Unification" here is geometric/entropic (shared edge diffusion parameter, heat-kernel weights). It is not an embedding into a simple Lie group. If the reconstructed gauge group genuinely factorizes as a product (sector factorization selector), there are no mixed generators playing the X/Y role. "Unify couplings" and "unify groups" therefore come apart.

10. Cosmological constant locality.

The cosmological constant problem is a graveyard of unified theories: local QFT estimates are enormous, and tiny observed Λ seems to demand absurd fine tuning.

OPH route: From null modular data, T_ab is reconstructed only up to ϕg_ab. Local consistency conditions and null focusing are blind to vacuum-energy shifts, so the Einstein equation is fixed only up to Λg_ab. Λ becomes a global "capacity" parameter of the static patch (tied to log dim H_tot). It is not a locally computable quantity. This resolves the conceptual tension: local microphysics *cannot* fix Λ by structural information-theoretic reasons.

11. UV infinities and nonrenormalizability.

Unified programs struggle to give sharp, finite microscopic definitions. Formal continuum structures, infinite entropies, and regularization dependence abound.

How OPH addresses it: The regulator premises explicitly require local patch algebras to be type-I and finite-dimensional, with a MaxEnt branch whose generator is quasi-local and obeys a Lieb-Robinson bound. So the fixed-cutoff UV branch is interacting in the ordinary finite-range sense, and the fundamental degrees of freedom are finite and live on the screen. What OPH does *not* claim is a unique microscopic UV completion: physical uniqueness is only modulo gauge or implementation hiding together with inert ancillary stabilization. The genuinely noncentral topological branch is also closed at fixed cutoff by the higher-gauge crossed-module collar theorem; the remaining burdens are the continuum BW lift and the refinement-limit transportable gauge branch.

12. Predictivity vs. parameter explosion.

Unified models often explode in parameters, sectors, or vacua, becoming hard to test directly because everything depends on choices.

How OPH addresses it: The framework compresses freedom into a "pixel area" (resolution) parameter and a total Hilbert space capacity (size) parameter, then derives structure from consistency (Lorentz, Einstein form, compact gauge group reconstruction, exact zeros, quantization patterns). The explicit MAR selector then picks the SM factors and sector factorization on the admissible class discussed in Section 6.2.

Structural pattern. OPH treats locality, Lorentz invariance, gauge symmetry, and gravity as consequences of consistency conditions among overlapping descriptions together with information-theoretic properties of states. Modular rigidity then supplies the familiar symmetry and dynamical structures.

Engineering deliverables. Certain problems can be stated as explicit closure tasks:

- Λ is structurally explained as a global capacity parameter and is not numerically predicted
- A full microphysical derivation of geometric modular action is still required

These are shared challenges across unification approaches. The framework provides an explicit map of where they live and what would resolve them.

---

4 Consensus, Defects, and Implementation Hiding

The OPH core admits an equivalent finite patch-net formulation. Local patch descriptions agree on overlaps, local repair rules remove mismatches, and the physical output is the schedule-independent normal form on the gauge quotient. This section records the fixed-point, defect, and gauge-quotient statements in that language.

The consensus lane also carries a simple law-selection model: once one equips candidate repair rules with a fitness functional, replicator dynamics gives a clean toy picture in which more successful reconciliation laws dominate a competing pool. That selection picture is useful for the interpretation of OPH, but it is not part of the recovered-core gravity/gauge theorem surface.

Theorem (Asynchronous confluence). On a finite patch net with local repair maps, if every accepted repair strictly decreases a finite inconsistency potential and the repair relation is locally confluent, then every initial configuration has a unique schedule-independent normal form.

Proof sketch. Strict decrease of the inconsistency potential rules out infinite repair sequences. Newman's lemma [5] then upgrades local confluence to global confluence. Therefore every initial state has a unique normal form, independent of the asynchronous update schedule.

Theorem (Cycle holonomy and higher-gauge defects). Pairwise overlap agreement is not enough for global consistency. On the abelian branch, a global solution exists iff the cycle holonomy vanishes on every loop. On the genuinely noncentral branch, weak gluing data define a crossed-module Cech class

$$q_\Sigma \in \check{H}^2(N_\Sigma, H_\Sigma \to G_\Sigma),$$

and strict global reconciliation exists iff $q_\Sigma = 0$. Nonzero q_Σ labels stable fixed-cutoff higher-gauge defects.

Proof sketch. In the abelian case, summing edge constraints around a loop telescopes, so vanishing of the loop sum is necessary and sufficient for path-independent reconstruction. In the crossed-module case, local rechartings are exactly crossed-module coboundaries, so they preserve the class and remove the defect precisely when the class is trivial.

Theorem (Gauge quotient, ancilla-stable physical uniqueness, and record closure). If the repair dynamics is gauge covariant, the normal-form map descends to the quotient by implementation hiding:

$$[s] \mapsto [\mathrm{nf}(s)].$$

Gauge-invariant observables are therefore uniquely fixed on the quotient, and inert ancillary refinement does not change the physical output. On the classical record layer, commuting closure operators on a finite lattice converge under any asynchronous fair schedule to the least common fixed point above the initial record state.

Proof sketch. Gauge covariance carries each repair sequence to a repair sequence on a gauge-related state, so uniqueness of normal forms forces equivariance of the normal-form map. Inert ancillary factors remain spectators under the lifted dynamics, so gauge-invariant observables are unchanged. For records, commutativity and idempotence make the closure order-independent and force convergence to the least common fixed point.

Extended derivation. The full consensus, gauge-quotient, and record-layer theorem package is developed in *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics* [2].

5 Particle-Spectrum Branch

The particle branch follows the same logical order as *Deriving the Particle Zoo from Observer Consistency* [4]. One first fixes the OPH basis and the compact gauge branch, then identifies the realized Standard Model quotient, and only then asks how transport-stable overlap data organize the observed particle families.

Structural particle theorem package. On the realized OPH branch, the structural carrier skeleton is fixed by the compact gauge and MAR-admissibility chain:

$$\frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}, \qquad N_g = 3, \qquad N_c = 3.$$

This determines the exact structural zero rows carried by the gauge and gravity sectors: the photon, gluons, and graviton are massless structural carriers on the realized branch.

Proof route. The proof route is the one developed in *A Conditional Reconstruction Program for Relativity and Standard Model Structure from Observer-Overlap Consistency: Observers Are All You Need (Compact)* [1] and *Deriving the Particle Zoo from Observer Consistency* [4]. Refinement-stable transportable edge sectors reconstruct a compact internal symmetry group under the stated categorical premises. MAR then selects the realized admissible branch, anomaly freedom fixes the hypercharge lattice, and the realized one-generation/one-Higgs package forces $N_g = 3$ and then $N_c = 3$. The particle branch builds on that structural skeleton and does not bypass it.

Sector	Current status	What is fixed here	Remaining boundary
Structural carriers	realized-branch structural theorem	exact $m_\gamma = m_g = m_{\mathrm{grav}} = 0$; realized quotient; $N_g = 3$; $N_c = 3$	theorem-grade structural exactness on the particle side
Electroweak bosons	calibration branch + exact sidecar	exact frozen-repair pair $M_W = 80.377$ GeV, $M_Z = 91.18797809193725$ GeV	exact only on the frozen authoritative repair surface; the public theorem rows remain the target-free electroweak outputs
Higgs/top stage	secondary quantitative branch + exact sidecars	exact Higgs inverse slice and exact top same-family witness	Higgs exactness is compare-only; top exactness is same-family-only
Quark family	continuation branch + exact witness	exact same-family quark sextet on the same ordered three-point family, with the ordered three-point quadratic readout closed on that scope	the selected quark sheet is the wrong branch; the next theorem-side object is `quark_d12_t1_value_law` on the emitted D12 mass ray, with `intrinsic_scale_law_D12` as the derived wrapper
Charged leptons	continuation branch + exact witness	exact same-family charged triple plus exact same-carrier centered readback, with the ordered three-point quadratic readout closed on that scope	same-family-only on the exact witness; theorem lane waits on $\widehat{C}_e^{\mathrm{cand}}$ and then $\mu_{\mathrm{phys}}(Y_e)$
Neutrinos	weighted-cycle continuation + exact adapter	exact compare-only two-parameter masses/splittings on the explicit positive selector segment	one reduced bridge-correction invariant C_ν remains open above the emitted proxy, so the exact adapter stays compare-only
Hadrons	nonperturbative continuation	stable-channel and readout architecture are defined	one production backend export bundle, then executed unquenching, runtime receipt, and production systematics

Status boundary. Structural carrier content and the realized Standard Model branch are theorem-bearing outputs. Electroweak calibration, the Higgs/top stage, quark rows, charged-lepton rows, neutrino rows, and hadron rows retain their stated calibration, continuation, or nonperturbative status.

Exact non-hadron hit surface. For quick orientation, the exact non-hadron lanes are:
The explicit values behind these exact-hit lanes are the exact frozen W/Z pair, the exact Higgs inverse slice, the exact same-family charged triple, the exact same-family quark sextet, and the exact compare-only neutrino mass triple with exact representative central splittings recorded in Ref. [4]'s exact-output surface.

Extended derivation. The structural-to-family route and the full per-sector status audit are developed in *Deriving the Particle Zoo from Observer Consistency* [4].

Lane	Exact output(s)	Exact OPH chain on the paper surface	Caveat
Structural carriers	$m_\gamma = m_g = m_{\text{grav}} = 0$	axioms → realized electromagnetic/color/dynamical-metric carrier skeleton → symmetry-protected zeros	theorem-grade structural exactness
Electroweak exact sidecar	exact frozen W/Z pair	axioms → $D7$ → $D10$ calibration chain → exact frozen repair pair	exact only on the frozen authoritative repair surface; compare-only beneath the target-free theorem
Higgs exact sidecar	exact Higgs inverse slice	axioms → $D10$ gauge core → $D11$ Jacobian → exact inverse slice	compare-only inverse slice
Charged exact witness	exact same-family charged triple	axioms → shared excitation dictionary → ordered charged carrier → exact centered readback → closed quadratic readout theorem → same-family exact witness	same-family-only; theorem lane waits on $\widehat{C}_e^{\text{cand}}$ and then $\mu_{\text{phys}}(Y_e)$
Quark exact witness	exact same-family quark sextet	axioms → shared excitation dictionary → D12 mass-side continuation → selected D12 sheet → closed quadratic readout theorem → same-family exact witness	same-family-only; the selected D12 sheet is the wrong CKM branch, and the next theorem-side object is `quark_d12_t1_value_law`, with `intrinsic_scale_law_D12` as the derived wrapper
Neutrino exact adapter	exact compare-only mass triple and representative central splittings	axioms → same-label scalar certificate → weighted-cycle branch → explicit positive selector segment → two-parameter exact adapter	compare-only; the theorem lane waits on the reduced bridge-correction invariant C_ν

6 Screen Microphysics, Records, and Observer Continuation

The screen-microphysics lane is the engineering side of the OPH suite. It asks whether some microscopic model might exist in principle and gives one explicit regulated reference architecture: a finite gauge-register or quantum-link style screen model in which local registers, patches, overlaps, edge observables, records, repair instruments, and observer-facing interfaces are all made concrete at fixed cutoff.

The status split on this lane has four pieces. First, the reference architecture itself is explicit. Second, the regulated patch-net embedding branch and the edge-sector thermal/Casimir branch are both sharpened, and each keeps one remaining closure pass visible. Third, the fixed-cutoff measurement/Born-rule package is closed. Fourth, the checkpoint/restoration observer package is closed. What remains open on this lane is not the existence of a workable finite regulator model, but the final closure of the patch-net embedding and compact heat-kernel lifts, global consensus analysis, refinement-stable universality, and the continuum/gravity lift.

Reference-architecture takeaway. The synthesis-level point is that OPH has a concrete screen-side habitat in which one can write local observables, overlap observables, record registers, and synchronization maps without pretending to identify a unique final UV completion.

Theorem (Fixed-cutoff record algebra and Born-Luders package). For each completed compare/write/verify slice of the regulated microphysics, the declared pointer and overlap-sector projectors generate a finite commutative record algebra

$$\mathcal{Z}_{\text{rec}}.$$

Every observer-accessible event E in that algebra has probability

$$\mathbb{P}(E) = \text{Tr}(\rho P_E),$$

and conditioning on E gives the operational post-measurement state

$$\rho|_E = \frac{P_E \rho P_E}{\mathrm{Tr}(\rho P_E)}.$$

Proof sketch. At fixed cutoff, the declared record and pointer projectors commute and the measurement surface factors through that finite classical algebra. Finite-dimensional projective measurement on a commuting projector algebra gives the Born trace and the Luders update directly.

Theorem (Checkpoint/restoration and observer backup). For an observer patch P_O, let the checkpoint data consist of the observer-facing record algebra, the accessible local state, the future update schedule, and the externally visible overlap interface data. Exact restoration reproduces the full future law of observer-accessible events, while an ε-accurate restoration changes that future law by at most ε in total variation.

Proof sketch. Exact restoration reinstates the same accessible state on the same observer-facing algebra with the same future channel sequence, so every later event probability agrees exactly. Approximate restoration is controlled by contractivity of trace distance under completely positive trace-preserving maps, which propagates the initial ε error bound to all later observer-accessible events.

Observer-facing conclusion. The theorem-level content is modest but sharp: OPH supplies a fixed-cutoff measurement interface, a fixed-cutoff checkpoint/restoration package, and an operational observer-identity criterion on observer-accessible event algebras. Stronger substrate-selection or strange-loop closure claims remain outside the recovered theorem package.

Extended derivation. The fixed-cutoff measurement, Born-rule, and checkpoint/restoration packages are developed in *Screen Microphysics, Patches, and Observer Synchronization in OPH* [3].

7 Conditional Worldsheet/String Branch

The synthesis paper carries the string/worldsheet branch at the same status as the dedicated string paper: a genuine theorem-level bridge from OPH edge sectors to the two-dimensional Yang-Mills partition function, followed by a conditional large-N_{edge} continuation to a worldsheet reorganization.

Theorem (Heat-kernel edge-sector bridge). Under the OPH axioms, the fixed-cutoff overlap-gauge realization, and the local-Gibbs/MaxEnt edge branch, the edge-sector weights satisfy

$$p_R(t) \propto d_R e^{-tC_2(R)}.$$

Consequently the edge partition function matches the standard two-dimensional Yang-Mills heat-kernel form.

Proof sketch. The fixed-cutoff collar package produces a sector decomposition on the edge algebra, and the local-Gibbs MaxEnt branch fixes the sector weights by the quadratic Casimir. Peter–Weyl decomposition [30] then identifies the resulting sum with the standard heat-kernel expression for two-dimensional Yang-Mills.

Continuation boundary. If a controlled large-N_{edge} regime exists, with N_{edge} distinct from the physical color rank $N_c = 3$, then the Gross-Taylor rewriting [29] gives a genus expansion and a worldsheet-like reorganization. Critical-superstring claims, worldsheet CFT closure, anomaly cancellation, and full massless-spectrum matching remain outside the recovered OPH core.

Extended derivation. This section records the synthesis-level statement; the longer technical discussion remains in the repository string fragment.

8 Cross-Lane Status and Open Directions

The state of the OPH suite is:

1. The fixed-cutoff collar, higher-gauge, and consensus theorem packages are part of the closed finite-regulator surface. The genuinely noncentral topological branch is not fixed-cutoff gap, and the quotient-level repair normal form is explicit on the finite patch-net surface.
2. The screen-microphysics lane has an explicit regulated reference architecture together with closed measurement/Born-rule and checkpoint/restoration packages. Its remaining open fronts are the final patch-net embedding pass, the compact heat-kernel lift, refinement-stable universality, global consensus closure, and the continuum/gravity lift.
3. Physical UV uniqueness is quotient-level and ancilla-stable: OPH fixes the physical branch modulo implementation hiding and inert ancillary refinement, not a unique microscopic representative.
4. The Lorentz/null-modular/Einstein chain remains a controlled refinement/scaling-limit branch. In the canonical premise ledger, the external burdens are the geometric modular branch, the scaling-limit scope clause, and the fixed-cap stationarity step.
5. Compact gauge reconstruction and the realized Standard Model quotient remain tied to the transportability, symmetric-braiding, and fiber-functor premises stated in the canonical technical-premise list.
6. The particle branch is structurally rich and sector-split: structural carriers and the realized Standard Model skeleton are fixed, electroweak closure is in hand, the Higgs/top stage is quantitatively emitted on its secondary branch, and the charged-lepton, quark, neutrino, and hadron closures sit on explicit continuation fronts.
7. The string branch is a continuation of the heat-kernel edge-sector theorem, not a promoted recovered-core result.

Companion papers. Detailed derivations appear in Refs. [1, 2, 3, 4].

9 Conclusion

Observer-Patch Holography is presented here as a single synthesis surface. The OPH basis, the conditional gravity and gauge branches, the consensus formulation, the particle-spectrum branch, the regulated screen-microphysics lane, the fixed-cutoff measurement and observer package, and the conditional worldsheet continuation appear on one common theorem and status surface.

The suite is uneven but concrete. OPH has a sharp fixed-cutoff local picture, a conditional recovered core for relativity and Standard-Model structure, real particle-side quantitative outputs, and a usable microscopic engineering lane. The main unfinished fronts are equally explicit: the continuum BW/geometric lift, the transportable-sector compact-gauge lift, the remaining microphysics closure passes, and the open matter-family and hadron closures. Refs. [1, 2, 3, 4] remain the depth surfaces for sector-by-sector proofs, calibration details, and continuation-level material.

A Candidate Microphysics Reference Architecture

The synthesis paper keeps the theorem-bearing observer and measurement package in the main text, but the concrete finite screen architecture belongs in an appendix. The working OPH reference architecture is a finite cellulation of the horizon screen with boundary-fixed patch algebras, boundary-fixed overlap algebras on collars, declared shared readout packets, local repair instruments, and record registers exposed on the observer-facing interface.

Theorem (Regulated patch-net embedding). On a fixed finite cellulation and finite patch cover, the regulated OPH microphysics realizes an explicit finite patch-net layer: patch vertices are genuine screen patches, overlap edges are genuine collars, packet fields are actual readouts of declared overlap observables, and allowed repair branches are concrete local channels on the same finite regulator.

Proof sketch. All patch and overlap algebras are finite-dimensional at fixed cutoff, and the declared packet fields are readouts of explicit overlap observables in those algebras. The candidate repair menu is implemented by local finite-support channels on the same neighborhoods. Therefore the abstract finite patch-net syntax is embedded into an explicit regulator object rather than introduced as a separate formal layer.

Theorem (Fixed-cutoff edge heat-kernel branch). On the thermalized overlap branch of the same regulated architecture, the stationary sector law is

$$\pi_\beta(\alpha) \propto d_\alpha e^{-\beta C_2(\alpha)},$$

and along refinement ladders converging to a compact-group Peter–Weyl decomposition [30] the same weights converge cylinderwise to the heat-kernel/Casimir law used in the OPH edge-sector branch.

Proof sketch. The regulated proposal kernel is chosen so that the weighted proposal symmetry $d_\alpha q_{\alpha\beta} = d_\beta q_{\beta\alpha}$ holds. Detailed balance with the Casimir acceptance rule then fixes the stationary distribution in finite dimension, and the compact-group heat-kernel law is the corresponding refinement lift.

Extended derivation. The full fixed-cutoff architecture, validation package, and theorem stacks are developed in *Screen Microphysics, Patches, and Observer Synchronization in OPH* [3].

B Interpretive Epilogue: Habitat and Strange-Loop Closure

This appendix is interpretive. It uses the OPH state-and-law habitat theorem and leaves the closure map on an observer-supporting sector open.

The OPH inputs used here are:

1. the patch net $P \mapsto A(P)$ of von Neumann algebras together with isotony and overlap restriction maps;
2. the global state on the inductive-limit algebra, whose restrictions witness a nonempty overlap-consistent local sector;
3. the finite Axiom-3 MaxEnt/refinement branch, which supplies a common finite family of gauge-invariant local constraint observables across cutoffs;
4. on the fixed-cutoff branch, the derived finite type-I presentations and compact boundary gluing groups used elsewhere in the OPH package.

The law slot is encoded by the finite expectation-value coordinates of the retained Axiom-3 constraint family. That choice furnishes a compact convex state-and-law habitat within OPH.

Theorem (OPH internal state-and-law habitat theorem). Fix finitely many screen patches $P_1, \ldots, P_N$ in the standard OPH setup. For each i, let

$$M_i := A(P_i), \qquad M_{ij} := A(P_i \cap P_j),$$

where M_i is the patch von Neumann algebra and $M_{ij} \subset M_i, M_j$ is the overlap algebra. Let

$$S_i := S(M_i) \subset M_i^*$$

be the full state space of M_i, endowed with the weak* topology $\sigma(M_i^*, M_i)$.

Choose a finite family of self-adjoint gauge-invariant local constraint observables

$$\mathcal{O} = \{O_1, \ldots, O_m\}$$

from the Axiom-3 MaxEnt branch, with each O_a supported in one of the chosen patches; write $i(a)$ for an index such that $O_a \in M_{i(a)}$.

Define the overlap-consistent state sector

$$X_{\text{ov}} := \left\{ (\omega_1, \ldots, \omega_N) \in \prod_{i=1}^{N} S_i : \omega_i|_{M_{ij}} = \omega_j|_{M_{ij}} \text{ for all } i, j \right\}.$$

Define the OPH law-coordinate map

$$c : X_{\text{ov}} \to \mathbb{R}^m, \qquad c(\omega_1, \ldots, \omega_N) := \big(\omega_{i(1)}(O_1), \ldots, \omega_{i(m)}(O_m)\big).$$

Let

$$X_{\mathcal{O}} := \left\{ (\omega_1, \ldots, \omega_N, \ell) \in \left(\prod_{i=1}^{N} S_i \right) \times \mathbb{R}^m : (\omega_1, \ldots, \omega_N) \in X_{\text{ov}},\ \ell = c(\omega_1, \ldots, \omega_N) \right\}.$$

Then:

1. the ambient product
$$E := \left(\prod_{i=1}^{N} M_i^*\right) \times \mathbb{R}^m$$
is a Banach space for any product norm;
2. $X_{\mathcal{O}}$ is nonempty and convex;
3. $X_{\mathcal{O}}$ is norm-closed in E;
4. $X_{\mathcal{O}}$ is compact in the product topology
$$\tau := \left(\prod_{i=1}^{N} \sigma(M_i^*, M_i)\right) \times \text{Euclidean topology on } \mathbb{R}^m;$$
5. therefore every τ-continuous self-map
$$\Phi : X_{\mathcal{O}} \to X_{\mathcal{O}}$$
has a fixed point.

Proof. Each M_i is a von Neumann algebra in the OPH patch net, hence a Banach space, so each dual M_i^* is Banach. A finite product of Banach spaces is Banach, hence so is E.

For each i, the state space

$$S_i = \{\omega \in M_i^* : \omega \geq 0,\ \omega(1) = 1\}$$

is convex and weak* compact by Banach–Alaoglu, because it is a weak* closed subset of the dual unit ball. For every pair (i, j), restriction along $M_{ij} \subset M_i$ and $M_{ij} \subset M_j$ defines affine weak* continuous maps

$$r_{ij} : S_i \to S(M_{ij}), \qquad r_{ji} : S_j \to S(M_{ij}).$$

Hence X_{ov} is an intersection of affine equalizer sets, so it is convex and τ-closed inside $\prod_i S_i$. Since $\prod_i S_i$ is compact and convex in the product weak* topology, X_{ov} is compact and convex as well.

Nonemptiness comes from the OPH global state on the inductive-limit algebra: its restrictions

$$(\omega|_{M_1}, \ldots, \omega|_{M_N})$$

belong to X_{ov} because the restrictions agree on every overlap by construction.

Each coordinate $\omega_{i(a)} \mapsto \omega_{i(a)}(O_a)$ is affine and weak* continuous, since evaluation at a fixed algebra element is weak* continuous. Thus c is affine and τ-continuous. The graph map

$$G : X_{\text{ov}} \to E, \qquad G(x) := (x, c(x)),$$

is affine and continuous, so its image $X_{\mathcal{O}} = G(X_{\text{ov}})$ is convex and τ-compact.

To see norm-closedness in E, let $(x_n, c(x_n)) \to (x, \ell)$ in norm. Norm convergence implies weak* convergence on each coordinate, so $x \in X_{\text{ov}}$ because X_{ov} is weak* closed. Since c is weak* continuous, $c(x_n) \to c(x)$, while norm convergence in $\mathbb{R}^m$ also gives $c(x_n) \to \ell$. Hence $\ell = c(x)$, so $(x, \ell) \in X_{\mathcal{O}}$.

Finally, $X_{\mathcal{O}}$ is a compact convex subset of the locally convex topological vector space E equipped with τ. Therefore every τ-continuous self-map $\Phi : X_{\mathcal{O}} \to X_{\mathcal{O}}$ has a fixed point by the Schauder–Tychonoff fixed-point theorem. □

Corollary (Fixed-cutoff OPH Brouwer corollary). On the fixed-cutoff realized presentation of OPH, each chosen patch algebra $M_i = A(P_i)$ is finite-dimensional. Then E is finite-dimensional, the weak* and norm topologies coincide on E, and $X_{\mathcal{O}}$ is a nonempty compact convex subset of a finite-dimensional real vector space. Hence every continuous self-map

$$\Phi : X_{\mathcal{O}} \to X_{\mathcal{O}}$$

has a fixed point by Brouwer. If, moreover, Φ is a contraction for some metric on $X_{\mathcal{O}}$, then the fixed point is unique and the iterates converge to it from every starting point by Banach's contraction theorem.

Proof. On that fixed-cutoff realized presentation, the OPH branch identifies each chosen patch algebra with a finite-dimensional type-I presentation modulo compact boundary redundancy, so each M_i and M_i^* is finite-dimensional. Hence E is finite-dimensional and all relevant locally convex topologies coincide. The habitat theorem above shows that $X_{\mathcal{O}}$ is nonempty, compact, and convex, so Brouwer gives a fixed point for every continuous self-map. The contraction statement is the usual Banach fixed-point theorem. ☐

The additional inputs needed for a strange-loop closure theorem are:

1. an actual OPH closure map Φ built from internal operations such as overlap repair, Axiom-3 MaxEnt reprojection, collar recoverability, and record-sector coarse-graining;
2. a proof that the stronger "observer-supporting" or "selected-world" criteria carve out a nonempty compact convex Φ-invariant subset of $X_{\mathcal{O}}$;
3. any uniqueness, contraction, or Lyapunov-type stability estimate for that closure map.

The habitat theorem supplies the Banach/compact-convex setting. A strange-loop closure theorem would additionally require the closure map and the stronger invariant-sector and stability inputs.

B.1 Timeless Causal Chain

The hypothesis is represented by the chain

information constraints → effective physics → complex chemistry and biology
→ observers with records → formal OPH reconstruction
→ OPH-compatible simulation substrate → information constraints.

In this appendix the chain is interpreted as a consistency loop. The discussion does not use it as a linear first-cause narrative. The Escher drawing-hands image is used only as a geometric analogy for this closure structure.

B.2 Why Reality Exists and Why It Has This Form

Under this hypothesis, existence is identified with nontrivial fixed points of the OPH consistency map. The observed form of reality is selected by stability requirements:

1. overlap-consistent sharable records,
2. recoverability under local information loss,

3. long-lived structured sectors that support observer formation,
4. internal reconstructibility of the governing rule set.

The appearance of intelligent observers enters the closure mechanism itself. Observer-capable worlds are the sectors that complete the loop.

B.3 Additional Problem Closures

Using the documentary node labels defined in Section 3.2.6, Phase I is D1–D5 together with D7–D9; D6 is the separate input-dependent corollary; and D12 collects phenomenological continuations. Ref. [3] proves fixed-cutoff theorem packages for the measurement/Born-rule interface and for checkpoint/restoration/backup. Table 14 records the status of the broader problem-closure claims discussed in the OPH synthesis surfaces.

Table 14: Problem-closure statements sorted by claim tier.

Problem domain	Status	OPH mechanism / scope note
Quantum gravity consistency	Phase I recovered core (conditional scaling-limit)	Conditional Lorentz and Jacobson-type Einstein branches from modular geometry, the null bridge, and fixed-cap stationarity.
UV completion / microscopic uniqueness	Mixed: fixed-cutoff theorem package / conditional scaling-limit branch	OPH gives a finite-range interacting fixed-cutoff branch, a unique physical UV branch only modulo OPH-stable equivalence, and a fixed-cutoff higher-gauge closure of the genuinely noncentral topological branch. Open items are a unique microscopic representative, the continuum BW lift behind T1/T3, and the refinement-limit transportable compact-gauge lift behind T2/T5–T7.
Measurement problem	Mixed: fixed-cutoff theorem package / interpretive synthesis	Ref. [3] proves a fixed-cutoff central-record / Born-Luders package for observer-accessible events. Broader philosophical closure claims remain interpretive here.
Cosmological principle and horizon homogeneity	Interpretive or organizational synthesis	MaxEnt state selection under symmetric constraints motivates isotropy and homogeneity stories. Those stories remain outside the independent Phase-I theorems.
Problem of time	Interpretive or organizational synthesis	Modular flow from the BW_{S^2} geometric branch provides a physical time parameter for each observer patch. It is an organizational principle rather than a separate recovered-core theorem.
Black-hole information paradox	Phenomenological continuation	Edge-center sectorization and recoverability suggest an internal resolution. The full black-hole-information branch remains outside Phase I.
Cosmological constant problem	Phase II input-dependent corollary	Λ is tied to global screen capacity once the screen-capacity identification is adopted. This sits outside the local recovered-core outputs.

Problem domain	Status	OPH mechanism / scope note
Magnetic monopole expectation from simple GUT groups	Phase I recovered core	The realized product-group branch $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)/\mathbb{Z}_6$ removes the standard simple-GUT symmetry-breaking monopole channel.
Proton stability and proton spin fraction	Mixed: Phase I / Phase III	No gauge-mediated proton decay follows from the realized product-group branch, which removes the standard simple-GUT proton-decay channel. Proton spin requires nonperturbative QCD completion and is not part of Phase I.
Dark matter phenomenology	Phase III deferred continuation	Modular-anomaly response laws remain conjectural beyond the recovered gravity chain.
Baryon asymmetry scale	Phase III deferred continuation	Suppression-counting estimates are not a substitute for a derived out-of-equilibrium mechanism.
Three generations and downstream flavor structure	Mixed: Phase I / Phase III	$N_g = 3$ is recovered on the realized MAR-admissible branch; later flavor constructions require extra ansatze.
String/worldsheet reorganization	Phase III deferred continuation	Requires a controlled large-N_{edge} regime distinct from the physical color rank $N_c = 3$. Any further lift to critical superstring structure remains a conjectural extension requiring extra worldsheet ingredients not derived in the OPH core.
Why anything exists / strange-loop closure	Interpretive epilogue only	Appendix B gives an OPH-internal state-and-law habitat theorem: overlap-consistent patch states together with finite Axiom-3 law coordinates form a theorem-usable nonempty compact convex sector, and fixed-cutoff branches admit a Brouwer version. Additional inputs are the OPH closure map on an invariant observer-supporting sector, together with uniqueness and stability.

C Observer Continuation and Backup

This appendix uses the fixed-cutoff checkpoint/restoration/error/identity theorem package proved in Ref. [3] and summarizes its algebraic interface in observer language. It does not define a strange-loop closure map on an invariant observer-supporting sector or establish uniqueness and stability for such a map.

C.1 Observer as Algebraic Pattern

In OPH, an observer is represented by

$$O = (P, \mathcal{A}(P), \rho, R),$$

where P is a screen patch, $\mathcal{A}(P)$ is its local algebra, ρ is the local state, and R is the record algebra. Records are implemented by approximately commuting projectors in overlap centers, so they are shareable without violating no-cloning constraints for generic quantum states.

Ref. [3] proves a fixed-cutoff observer-facing measurement interface: a finite central record algebra, Born probabilities for its event projectors, and Luders conditioning on that same commuting record algebra. This appendix uses that fixed-cutoff measurement package in observer language.

C.2 Markov Collar Factorization

For a collar tripartition A-B-D with conditional mutual information $I(A : D|B) \leq \varepsilon$, the exact Markov limit gives

$$\rho_{ABD} = \bigoplus_{\alpha} p_\alpha \, \rho^{(\alpha)}_{Ab^\alpha_L} \otimes \rho^{(\alpha)}_{b^\alpha_R D}.$$

The sector label α is classical center data. This decomposition is the mathematical basis for extracting an interior observer state with a controlled boundary interface.

C.3 Checkpoint and Restoration Map

A checkpoint is the tuple

$$\mathcal{C} = (R, \alpha, \rho^{(\alpha)}_{\text{int}}),$$

with $\rho^{(\alpha)}_{\text{int}}$ the interior conditional state. Given a compatible target environment state $\sigma^{(\alpha)}_{\text{env}}$, a restored state is

$$\rho^{(\alpha)}_{\text{new}} = \rho^{(\alpha)}_{\text{int}} \otimes \sigma^{(\alpha)}_{\text{env}}.$$

For approximate Markov collars, recovery is controlled by the standard bound

$$\|\rho_{ABD} - (\text{id}_A \otimes \mathcal{R}_{B\to BD})(\rho_{AB})\|_1 \leq 2\sqrt{\ln 2\,\varepsilon}.$$

This gives quantitative trace-distance control on the recovered state under finite conditional mutual information (finite collar error). Interpreting that control as observer continuation requires additional modeling assumptions beyond the bound itself.

The fixed-cutoff microphysics claim is stronger than the bare recoverability estimate written above. In Ref. [3], exact checkpoint restoration preserves the full future law of observer-accessible events, while an ε-accurate restoration changes that future law by at most ε in total variation. The stronger strange-loop closure story remains open.

C.4 Physical Meaning

At fixed cutoff, Ref. [3] proves a checkpoint/restoration theorem and backup corollary for observer-accessible event algebras. This appendix states the observer-facing meaning of that theorem and its algebraic prerequisites. Open items are the stronger substrate-selection, strange-loop closure, uniqueness, and stability package.

References

[1] B. Müller and P. Nguyen, *A Conditional Reconstruction Program for Relativity and Standard Model Structure from Observer-Overlap Consistency: Observers Are All You Need (Compact)*, 2026. Available at `https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/recovering_relativity_and_standard_model_structure_from_observer_overlap_consistency_compact.pdf`.

[2] B. Müller, *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics*, 2026. Available at `https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/reality_as_consensus_protocol.pdf`.

[3] B. Müller, *Screen Microphysics, Patches, and Observer Synchronization in OPH*, 2026. Available at `https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/screen_microphysics_and_observer_synchronization.pdf`.

[4] B. Müller, *Deriving the Particle Zoo from Observer Consistency*, 2026. Available at `https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/deriving_the_particle_zoo_from_observer_consistency.pdf`.

[5] M. H. A. Newman, "On theories with a combinatorial definition of 'equivalence'," *Ann. of Math. (2)* **43** (1942), 223–243.

[6] E. H. Lieb and D. W. Robinson, "The finite group velocity of quantum spin systems," *Commun. Math. Phys.* **28** (1972), 251–257.

[7] J. J. Bisognano and E. H. Wichmann, "On the duality condition for a Hermitian scalar field," *J. Math. Phys.* **16** (1975), 985–1007.

[8] J. J. Bisognano and E. H. Wichmann, "On the duality condition for quantum fields," *J. Math. Phys.* **17** (1976), 303–321.

[9] R. Brunetti, D. Guido, and R. Longo, "Modular Structure and Duality in Conformal Quantum Field Theory," *Commun. Math. Phys.* **156** (1993), 201–219, arXiv:funct-an/9302008.

[10] H.-W. Wiesbrock, "Half-Sided modular inclusions of von-Neumann-Algebras," *Commun. Math. Phys.* **157** (1993), 83–92.

[11] W. G. Unruh, "Notes on black-hole evaporation," *Phys. Rev. D* **14** (1976), 870–892.

[12] T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," *Phys. Rev. Lett.* **75** (1995), 1260–1263, arXiv:gr-qc/9504004.

[13] T. Jacobson, "Entanglement equilibrium and the Einstein equation," *Phys. Rev. Lett.* **116** (2016), 201101, arXiv:1505.04753.

[14] R. Bousso, Z. Fisher, J. Koeller, S. Leichenauer, and A. C. Wall, "Proof of the quantum null energy condition," *Phys. Rev. D* **93** (2016), 024017, arXiv:1509.02542.

[15] E. H. Lieb and M. B. Ruskai, "Proof of the strong subadditivity of quantum-mechanical entropy," *J. Math. Phys.* **14** (1973), 1938–1941.

[16] D. Petz, "Sufficient subalgebras and the relative entropy of states of a von Neumann algebra," *Commun. Math. Phys.* **105** (1986), 123–131.

[17] D. Petz, "Sufficiency of channels over von Neumann algebras," *Quart. J. Math.* **39** (1988), 97–108.

[18] O. Fawzi and R. Renner, "Quantum conditional mutual information and approximate Markov chains," *Commun. Math. Phys.* **340** (2015), 575–611, arXiv:1410.0664.

[19] P. Hayden, R. Jozsa, D. Petz, and A. Winter, "Structure of states which satisfy strong subadditivity of quantum entropy with equality," *Commun. Math. Phys.* **246** (2004), 359–374.

[20] S. Doplicher and J. E. Roberts, "A new duality theory for compact groups," *Invent. Math.* **98** (1989), 157–218.

[21] S. Doplicher and J. E. Roberts, "Why there is a field algebra with a compact gauge group describing the superselection structure in particle physics," *Commun. Math. Phys.* **131** (1990), 51–107.

[22] T. Tannaka, "Über den Dualitätssatz der nichtkommutativen topologischen Gruppen," *Tohoku Math. J.* **45** (1938), 1–12.

[23] M. G. Krein, "A principle of duality for a bicompact group and a square block algebra," *Dokl. Akad. Nauk SSSR* **69** (1949), 725–728.

[24] H. Georgi and S. L. Glashow, "Unity of all elementary-particle forces," *Phys. Rev. Lett.* **32** (1974), 438–441.

[25] S. L. Glashow, J. Iliopoulos, and L. Maiani, "Weak interactions with lepton-hadron symmetry," *Phys. Rev. D* **2** (1970), 1285–1292.

[26] E. Witten, "An SU(2) anomaly," *Phys. Lett. B* **117** (1982), 324–328.

[27] S. Dimopoulos, S. Raby, and F. Wilczek, "Supersymmetry and the scale of unification," *Phys. Rev. D* **24** (1981), 1681–1683.

[28] U. Amaldi, W. de Boer, and H. Fürstenau, "Comparison of grand unified theories with electroweak and strong coupling constants measured at LEP," *Phys. Lett. B* **260** (1991), 447–455.

[29] D. J. Gross and W. Taylor, "Two-dimensional QCD is a string theory," *Nucl. Phys. B* **400** (1993), 181–208, arXiv:hep-th/9301068.

[30] F. Peter and H. Weyl, "Die Vollständigkeit der primitiven Darstellungen einer geschlossenen kontinuierlichen Gruppe," *Math. Ann.* **97** (1927), 737–755.

[31] A. Bullivant, M. Calçada, Z. Kádár, P. Martin, and J. Faria Martins, "Topological phases from higher gauge symmetry in 3+1D," *Phys. Rev. B* **95** (2017), 155118, arXiv:1606.06639.

[32] S. Chandrasekharan and U.-J. Wiese, "Quantum link models: A discrete approach to gauge theories," *Nucl. Phys. B* **492** (1997), 455–471, arXiv:hep-lat/9609042.

[33] W. Donnelly and A. C. Wall, "Entanglement entropy of electromagnetic edge modes," *Phys. Rev. Lett.* **114** (2015), 111603, arXiv:1412.1895.

[34] F. Pastawski, B. Yoshida, D. Harlow, and J. Preskill, "Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence," *JHEP* **06** (2015), 149, arXiv:1503.06237.

[35] M. A. Levin and X.-G. Wen, "String-net condensation: A physical mechanism for topological phases," *Phys. Rev. B* **71** (2005), 045110, arXiv:cond-mat/0404617.

[36] D. Laghi, G. Carullo, J. Veitch, and W. Del Pozzo, "Quantum black hole spectroscopy: probing the quantum nature of the black hole area using LIGO-Virgo ringdown detections," *Class. Quantum Grav.* **38** (2021), 095005, arXiv:2011.03816.

[37] Particle Data Group, "Review of Particle Physics," *Phys. Rev. D* **110** (2024), 030001. Available at `https://pdg.lbl.gov/2024/download/db2024.pdf`.

[38] Particle Data Group, "Electroweak Model and Constraints on New Physics," Available at `https://pdg.lbl.gov/2024/reviews/rpp2024-rev-standard-model.pdf`.